Progressive Wealth Taxation in the Combined Flow–Stock Framework: Optimal Inequality Reduction under Gini and Top-Share Targets

Anders G. Frøseth*

July 8, 2026

Abstract

We characterise optimal progressive wealth taxation in coexistence with conventional flow taxes on capital ownership under two redistributive criteria: the Gini coefficient and the Saez–Zucman top share. Working within the Fokker–Planck framework on log-wealth, we extend the (C1)–(C3) neutrality class of Frøseth (2026d) to a combined flow-tax retention factor $k = (1-\tau_c)(1-\tau_d)$ paired with a state-dependent wealth-tax rate $r(x)$, and treat three analytically tractable sub-classes: smooth log-progressive (the small-rate linearisation of the Heathcote–Storesletten–Violante parametrisation), single-bracket, and two-bracket. The smooth and single-bracket sub-classes preserve criterion equivalence: both criteria collapse to monotone functions of a single shape parameter and the Gini- and top-share-targeted optimal schedules coincide. The two-bracket sub-class — the form actually implemented in Norway since 2023 — has a three-dimensional inequality- relevant policy space, and the criterion choice becomes binding: at a fixed Gini level the top-1% share ranges over roughly an order of magnitude (a factor of seven to eight in the Norwegian calibration) across the feasible iso-Gini manifold. At calibrated Norwegian regime parameters the wealth-tax channel alone is below the steady-state confinement threshold; jointly with the demographic turnover of Frøseth (2026f, §3.2) and the emigration-contagion channel of Frøseth (2026a) the framework recovers the empirical Pareto-tail exponent from the Kapital 400 panel. The decomposition reveals that the current wealth tax contributes about a sixth of the structural compression beyond the no-tax reference, against five-sixths from the flow-tax channel; and that, within the homogeneous-returns class, the structural effect of the wealth tax is replicable by a moderate flow-tax increase. Outside that class, on the dormant-wealth base and under heterogeneous returns, the two taxes have structurally distinct incidence (Domar and Musgrave, 1944; Stiglitz, 1969; Frøseth, 2026c; Guvenen et al., 2023) and a heterogeneous-returns extension to the present framework is identified as the natural next step.

JEL Classification: D31, H21, H23, H24, C61.

Keywords: Wealth tax, progressive taxation, Fokker–Planck equation, Gini coefficient, top share, optimal taxation, Saez–Zucman, bracket schedule, Ornstein–Uhlenbeck process.

*Independent Researcher. E-mail: indrefjorden@pm.me.

Contents

1 Introduction

1.1 The policy question

Two strands of the contemporary tax-policy literature converge on the same instrument from different directions. The first, in the tradition of Piketty (2014); Piketty and Zucman (2014); Saez and Zucman (2019), argues that taxation of the wealth *stock* is the central redistributive lever in advanced economies, on the grounds that capital accumulation outpaces labour income at the top of the distribution. A progressive schedule on the stock—taxing larger fortunes at higher rates—is the policy instrument that follows from this view. The second, in the tradition of dual-income taxation (Sørensen, 2005; Alstadsæter et al., 2019), treats capital ownership as a *system* of taxes: a corporate tax on profits, a dividend or capital-gains tax on distributions, and (in some jurisdictions) a wealth tax on net assets, the whole calibrated to balance horizontal equity, lock-in, and avoidance margins. From this second view the wealth tax is one component of the ownership-tax system, and the design question is how it should interact with the others.

A policy-relevant theory must speak to both views. The distributional question—how progressive should the wealth tax be, judged against a redistributive criterion—requires a schedule class flexible enough to admit non-proportional rates on the stock, and an inequality measure (Gini, top-share) that the political debate actually centres on. The systemic question—how the wealth tax interacts with corporate and dividend taxes at the household level—requires the post-tax wealth process to carry the (C1)–(C3) decomposition of Frøseth (2026d), in which the flow-tax channels reduce to a single retention factor $k = (1 - \tau_c)(1 - \tau_d)$ acting on drift and diffusion. The present paper combines the two in a single optimal-design problem.

1.2 Position in the FP design programme

The Fokker–Planck framework for wealth taxation now consists of four papers we draw on directly. Frøseth (2026g) formalises the proportional wealth tax as a uniform drift-shift on log-wealth. Frøseth (2026f) extends the framework from drift-shift to *drift-design*: progressive schedules generate state-dependent drift, the Pareto-tailed steady state of the Bouchaud–Mézard model is replaced by an Ornstein–Uhlenbeck confinement, and the Gini coefficient becomes a closed-form function of the progressivity parameter. The criterion in that paper is the Gini, the schedule class is stock-only, and the analytical core is the OU steady state. Frøseth (2026d) extends the framework from stock-only to *flow + stock*: corporate, dividend, and capital-income taxes enter alongside the wealth tax, the (C1)–(C3) class preserves portfolio neutrality, and the post-tax SDE becomes a drift-shift-and-rescale of the pre-tax SDE. The criterion is implicit (neutrality), the schedule is proportional in the wealth-tax rate, and the analytical core is the neutrality-preserving structure. Frøseth (2026b) combines the two by characterising the optimal proportional wealth tax *within* the (C1)–(C3) flow-tax envelope under two information-theoretic and transport-geometric criteria (the JKO free-energy gap and the 2-Wasserstein distance from no-tax). The criterion is distortion-minimising, the schedule is proportional, and the analytical core is a phase-transition partition of the regime axis.

The natural completion of this two-by-two diagonal—progressive schedule class, redistributive criterion, in coexistence with flow taxes—is the present paper. Concretely:

	stock-only	*flow + stock*
proportional	Frøseth (2026g)	Frøseth (2026d)
progressive	Frøseth (2026f) (Gini target)	*this paper*

The criterion-family axis runs orthogonally: Frøseth (2026b) treats the proportional sub-class under JKO/W_2; the present paper treats the progressive sub-class under Gini/top-share.

1.3 Schedule class and criteria

We work in a two-channel design space. The flow-tax channel is parametrised by the corporate–dividend retention $k = (1 - \tau_c)(1 - \tau_d) \in (0, 1]$, identical to its role in Frøseth (2026d,b). The wealth-tax channel is a function $r(x)$, the per-unit rate on raw wealth x, allowed to depend on the wealth level. Three analytically tractable cases are studied. The *smooth log-progressive* case

$$r(x) \;=\; r_0 + \beta \log(x/x_0),$$

the leading-order linearisation of the Heathcote–Storesletten–Violante progressivity parametrisation in the small-rate regime, gives a linear restoring force on log-wealth, with the post-tax process an Ornstein–Uhlenbeck process with parameters (k, β). The *single-bracket* case

$$r(x) \;=\; r \cdot \mathbf{1}\{x \geq X_0\}$$

gives a piecewise-constant drift on log-wealth, with steady-state density an asymmetric Laplace distribution centred at $\log X_0$ (equivalently, a piecewise power-law density on raw wealth with Pareto-tail exponent $\alpha = 2(r/k - m_0)/\sigma^2$). The *two-bracket* case

$$r(x) \;=\; r_1 \mathbf{1}\{X_0 \leq x < X_1\} + r_2 \mathbf{1}\{x \geq X_1\}\,,$$

the form actually implemented in Norway since 2023 and several Swiss cantons, generalises this to a three-regime piecewise-exponential density with two distinct upper-tail exponents.

The two redistributive criteria evaluated on each schedule class are the Gini coefficient and the top share S_q, $q \in \{99, 99.9, \ldots\}$. The smooth and single-bracket classes are inequality-relevant only through a one-dimensional shape parameter (the structural ratio β/k and the Pareto exponent α respectively), so both criteria collapse to monotone functions of that parameter and pick the same optimal schedule up to reparametrisation. The two-bracket class has a three-dimensional inequality-relevant policy space, and the two criteria pick distinct optima within it: the choice of criterion genuinely affects *where* the optimal redistributive action sits.

1.4 Contributions and roadmap

The paper makes the following contributions.

1. A combined flow + progressive design problem on the (C1)–(C3) class, posed for the first

time. The post-tax SDE on log-wealth $Y = \log X$ is shown (Section 2) to take the form $\mathrm{d}Y = [k\, m_0 - r(e^Y)]\,\mathrm{d}t + \sqrt{k}\,\sigma\,\mathrm{d}W$, cleanly separating the flow-tax retention from the progressive wealth-tax channel.

2. Closed-form Gini and top-share expressions for the smooth log-progressive case (Section 3) reducing to the structural ratio $\rho_p = \beta/k$, with the optimum on the (C1)–(C3) class characterised by a single equation in ρ_p that recovers the Frøseth (2026f) result on the boundary $k = 1$. Section 3 grounds the choice of functional form as the small-rate linearisation of the Heathcote–Storesletten–Violante progressivity parametrisation, and develops the suppression of the Stiglitz (1985)-style price feedback channel under (C1)–(C3) neutrality.

3. Closed-form analysis of the bracket case (Section 4) via the asymmetric Laplace steady-state density on log-wealth, with explicit Gini and top-share expressions in elementary functions of the Pareto exponent α. The single-bracket sub-class preserves the criterion-equivalence baseline; the two-bracket sub-class breaks it. We exhibit numerically a parameter region where the same Gini level admits a roughly near-order-of-magnitude range of top-1% shares.

4. A Norwegian calibration (Section 6) showing that the actual two-bracket schedule (since 2023, with $X_0 \approx 1.7$ M NOK at $r_1 \approx 1.0\%$ and $X_1 \approx 20$ M NOK at $r_2 \approx 1.1\%$) sits below the steady-state confinement threshold under the wealth-tax-only mechanism. The empirical confinement is supplied jointly by the Gabaix-style demographic turnover of Frøseth (2026f, §3.2) and the emigration-contagion channel of Frøseth (2026a); we derive the joint Pareto-exponent formula and show that the criterion divergence of contribution (3) is robust across these confinement mechanisms.

5. A decomposition of the Norwegian wealth tax's structural redistributive contribution (Section 6.6). The current schedule contributes about 17% of the upper-tail compression relative to the no-tax baseline; the flow-tax channel contributes the remaining 83%. Within the homogeneous-returns class the structural effect of the wealth tax is in principle replicable by raising the dividend tax by approximately seven percentage points. Outside that class — on the dormant-wealth base, and under heterogeneous returns (Domar and Musgrave, 1944; Stiglitz, 1969; Frøseth, 2026c; Guvenen et al., 2023) — flow and stock taxes have structurally distinct incidence and the substitutability fails. A heterogeneous-returns extension to the design problem is identified as the natural next step.

Section 2 introduces the schedule class, the post-tax SDE, and the two criteria. Section 3 treats the smooth log-progressive strand. Section 4 treats the bracket strand. Section 5 bridges the framework to actual schedules through the assessment-fraction rescaling and the discrete-time approximation. Section 6 carries out the Norwegian calibration and integrates the joint analysis with demographic turnover and emigration. Section 7 positions the contribution within the broader programme. Section 8 closes with extensions and limitations.

2 Setup

This section introduces the wealth-process model, the combined flow + progressive schedule class on which the paper operates, and the two redistributive criteria evaluated on it. The exposition recapitulates only those elements of Frøseth (2026g,d,f) that are needed for the present analysis.

2.1 Wealth dynamics under no taxation

A representative investor holds a portfolio of risky and risk-free assets and faces the geometric Brownian motion

$$\frac{\mathrm{d}X_t}{X_t} \;=\; \mu\,\mathrm{d}t + \sigma\,\mathrm{d}W_t\,, \tag{1}$$

where X_t is wealth, μ is the gross expected return, σ is the portfolio volatility, and W is a standard Brownian motion. As in Frøseth (2026g), we work in log-wealth coordinates $Y_t = \log X_t$, where the process becomes

$$\mathrm{d}Y_t \;=\; m_0\,\mathrm{d}t + \sigma\,\mathrm{d}W_t, \qquad m_0 \;=\; \mu - \sigma^2/2\,, \tag{2}$$

with m_0 the geometric-mean drift. The associated Fokker–Planck equation for the density $\pi(y,t)$ of log-wealth is

$$\partial_t\pi \;=\; -\,m_0\,\partial_y\pi + \tfrac{1}{2}\sigma^2\,\partial_y^2\pi\,. \tag{3}$$

A heterogeneous-investor extension with idiosyncratic returns gives the same FP equation for the cross-sectional density, following the Bouchaud–Mézard tradition (Bouchaud and Mézard, 2000); we suppress the cross-sectional formalism here as it is not needed.

2.2 The (C1)–(C3) flow-tax envelope

We restrict attention throughout to the *(C1)–(C3) neutrality class* of Frøseth (2026d). The three conditions are:

(C1) The capital-income tax rate equals the corporate tax rate: $\tau_k = \tau_c$.

(C2) The shielding rate equals the risk-free rate: $r_s = r_f$.

(C3) The wealth-tax assessment is uniform across asset classes: $\alpha_i = \alpha$ for all i.

Under (C1)–(C3), the corporate–dividend–capital-income tax package collapses (Frøseth, 2026d, Theorem 1) to a single multiplicative *retention factor*

$$k \;=\; (1-\tau_c)(1-\tau_d) \;\in\; (0,1]\,, \tag{4}$$

acting uniformly on excess returns. The wealth-tax channel remains as a separate state-coupled component.

The post-tax wealth dynamics under (C1)–(C3) plus a wealth-tax schedule with rate $r(x)$ on raw wealth (i.e. the household pays $r(x)\,x\,\mathrm{d}t$ in wealth tax over the interval $\mathrm{d}t$) take the form

$$\frac{\mathrm{d}X_t}{X_t} \;=\; \left[k\mu - r(X_t)\right]\mathrm{d}t + \sqrt{k}\,\sigma\,\mathrm{d}W_t\,. \tag{5}$$

Equivalently, on log-wealth $Y_t = \log X_t$,

$$\mathrm{d}Y_t \;=\; \left[k\,m_0 - r(e^{Y_t})\right]\mathrm{d}t + \sqrt{k}\,\sigma\,\mathrm{d}W_t\,, \tag{6}$$

with $m_0 = \mu - \sigma^2/2$ the no-tax geometric-mean drift. The flow-tax channel enters *linearly* through k (rescaling drift uniformly by k and diffusion uniformly by k); the progressive wealth-tax channel enters through the state-dependent drift correction $r(e^Y)$. The two channels do not interact, in the sense of additive separability (Frøseth, 2026d, §7).

2.3 The combined flow + progressive schedule class

The schedule space studied in this paper is

$$\mathcal{S}_{\text{flow+prog}} \;=\; \big\{(k, r(\cdot)) \;:\; k \in (0,1],\; r(\cdot) : (0,\infty) \to [0, r_{\max}],\; r \text{ progressive in } x\big\}\,, \tag{7}$$

where "progressive in x" means non-decreasing as a function of raw wealth.[1]

Two analytically tractable subclasses receive most of our attention. The *smooth log-progressive* subclass

$$r(x) \;=\; r_0 + \beta\log(x/x_0)\,, \qquad (r_0, \beta) \in [0,\infty)^2, \tag{8}$$

parametrised by the floor rate $r_0 \geq 0$, the progressivity slope $\beta \geq 0$, and a normalising scale $x_0 > 0$, gives a state-coupled drift that is linear in $\log x$. As we show in Section 3, this is the unique subclass on which the post-tax process on Y is an Ornstein–Uhlenbeck process and the steady-state distribution of raw wealth is log-normal. The *bracket* subclass

$$r(x) \;=\; r\cdot\mathbf{1}\{x \geq X_0\}\,, \qquad (X_0, r) \in (0,\infty)\times[0, r_{\max}], \tag{9}$$

gives a piecewise-constant drift on log-wealth—zero below the threshold $\log X_0$, a downward shift $-r$ above—and is the form actually implemented in jurisdictions such as Norway, Switzerland, and (historically) Spain, France, and several other European countries. We extend the bracket subclass to two brackets where the additional generality is needed:

$$r(x) \;=\; r_1\cdot\mathbf{1}\{x \in [X_0, X_1)\} + r_2\cdot\mathbf{1}\{x \geq X_1\}\,, \quad r_2 > r_1 \geq 0\,. \tag{10}$$

The number of inequality-relevant policy levers within each subclass depends on how many of the schedule parameters can be reduced by the regime structure. As we derive in Section 3 the smooth subclass collapses to the single ratio $\rho_p = \beta/k$ at fixed regime, and as we derive in Section 4 the single-bracket subclass collapses to the Pareto exponent $\alpha = 2(r/k - m_0)/\sigma^2$ alone. The two-bracket subclass retains a three-dimensional inequality-relevant policy space at fixed

[1] We allow r to be discontinuous, as in the bracket case below; piecewise-monotone-non-decreasing is the operative regularity. Continuous regulation of r at discontinuities is not needed for any of our results.

regime: the substantive divergence content of the paper appears precisely at this dimensionality. The proportional sub-class of Frøseth (2026b) corresponds to $\beta = 0$ in the smooth subclass and to $X_0 = 0$ in the bracket subclass.

2.4 Two redistributive criteria

We evaluate every schedule in $\mathcal{S}_{\text{flow+prog}}$ against two redistributive criteria, both standard in the empirical inequality literature.

Gini coefficient. For a distribution of raw wealth with cumulative distribution function F and finite mean $\bar{X} = \int x \, \mathrm{d}F(x)$, the Gini coefficient is

$$\mathrm{Gini}[F] \;=\; \frac{1}{2\bar{X}} \iint |x - x'| \, \mathrm{d}F(x) \, \mathrm{d}F(x') \;=\; 2\int_0^1 [u - L(u)] \, \mathrm{d}u \,, \tag{11}$$

where $L : [0,1] \to [0,1]$ is the Lorenz curve. For log-normal raw wealth with log-variance σ_Y^2,

$$\mathrm{Gini} \;=\; 2\Phi\Big(\sigma_Y/\sqrt{2}\Big) - 1 \;=\; \mathrm{erf}(\sigma_Y/2) \,, \tag{12}$$

a closed form we use throughout the smooth strand.

Top share. For $q \in (0,1)$, the top share is the fraction of total wealth held by the top $(1-q)$ percentile:

$$S_q[F] \;=\; \frac{\int_{F^{-1}(q)}^{\infty} x \, \mathrm{d}F(x)}{\int_0^{\infty} x \, \mathrm{d}F(x)} \,. \tag{13}$$

For log-normal raw wealth with parameters (μ_Y, σ_Y^2),

$$S_q \;=\; 1 - \Phi\Big(\Phi^{-1}(q) - \sigma_Y\Big) \,, \tag{14}$$

also a closed form.

The two criteria are policy-relevant in different traditions: the Gini is the standard summary statistic across the distribution, while the top share is the focal measure in the Saez–Zucman tradition (Saez and Zucman, 2019). On log-normal distributions both are monotone in the same shape parameter σ_Y, so the optimal schedule is criterion-invariant. On bracket schedules the post-tax density is piecewise-Gaussian rather than log-normal, and the two criteria can disagree about *where* on the joint design plane the optimal redistributive action lies.

3 The smooth log-progressive strand

This section treats the smooth log-progressive subclass (8), $r(x) = r_0 + \beta \log(x/x_0)$. The post-tax process on log-wealth is an Ornstein–Uhlenbeck process, the steady-state distribution of raw wealth is log-normal, and both redistributive criteria admit closed forms on the two-dimensional design plane (k, β). The substantive content of the section is threefold: a small-rate-linearisation argument that grounds the smooth functional form in the standard public-finance progressivity parametrisation, a structural-ratio reduction that collapses the inequality dependence to a single combination β/k, and a comparison theorem stating that the Gini-optimal and top-share-optimal schedules coincide on the smooth class.

3.1 Functional-form motivation: the smooth schedule as a small-rate linearisation of HSV progressivity

The smooth log-progressive functional form $r(x) = r_0 + \beta \log(x/x_0)$ is not arbitrary. It is the leading-order Taylor expansion of the long-standing *constant-residual-progression* parametrisation in public finance — introduced by Feldstein (1969) and Persson (1983) in static optimal-taxation settings, brought into dynamic heterogeneous-agent macro by Bénabou (2002), and most recently developed in the analytical framework of Heathcote et al. (2017), where it has become the standard parametrisation. We refer to it as the Heathcote–Storesletten–Violante (HSV) form throughout.

The HSV form expresses the tax liability at base level y as $T_{\text{HSV}}(y) = y - \lambda y^{1-\tau}$ (Heathcote et al., 2017, equation (1)), with post-tax disposable level $\tilde{y} = \lambda y^{1-\tau}$ and average rate

$$\bar{t}_{\text{HSV}}(y) \;=\; 1 - \lambda\, y^{-\tau}\,, \tag{15}$$

where $\tau \in [0, 1]$ is a progressivity index and $\lambda > 0$ a level parameter. The functional form is the unique two-parameter family for which the *retention* rate $1 - \bar{t}_{\text{HSV}}(y)$ is log-linear in $\log y$, so that $\log(1 - \bar{t}_{\text{HSV}}(y)) = \log \lambda - \tau \log y$. The parameter $1 - \tau$ has a clean economic reading as the elasticity of post-tax to pre-tax base (Heathcote et al., 2017, equation (2)); the ratio of marginal to average retention rates is also $1 - \tau$, which is why τ is a natural index of *residual* progressivity. The case $\tau = 0$ recovers a proportional schedule ($\bar{t} = 1 - \lambda$); the case $\tau = 1$ delivers full progressivity ($\bar{t} \to 1$ as $y \to \infty$). HSV applied to wealth, $\bar{t}_{\text{HSV}}(x) = 1 - \lambda x^{-\tau}$, gives a progressive wealth-tax schedule with retention rate log-linear in log-wealth.[2]

The smooth log-progressive schedule (8) arises as the small-rate Taylor expansion of (15).

Proposition 1 (HSV linearisation)**.** *On any bounded log-wealth interval $\log x \in [\log x_0 - L, \log x_0 + L]$, the HSV-on-wealth schedule (15) and the smooth log-progressive schedule (8) agree to leading order in $\tau \log(x/x_0)$ when the parameters are linked by*

$$r_0 \;=\; 1 - \lambda x_0^{-\tau}\,, \qquad \beta \;=\; \lambda \tau\, x_0^{-\tau}\,, \tag{16}$$

with remainder bound

$$\left|\bar{t}_{HSV}(x) - r(x)\right| \;\le\; \tfrac{1}{2}\,\lambda x_0^{-\tau}\big(\tau \log(x/x_0)\big)^2 \;=\; \tfrac{1}{2}\,(1 - r_0)\big(\tau \log(x/x_0)\big)^2\,. \tag{17}$$

[2] In Bénabou's notation (Bénabou, 2002, eq. 6) the same family appears as $\tilde{y}_i = \bar{y}^{1-\phi}\, y_i^{\phi}$, where ϕ is the elasticity of disposable to market base. The mapping is $\phi = 1 - \tau$ and $\lambda = \bar{y}^{\,\tau}$, i.e. Bénabou absorbs the normalising scale into a balanced-budget condition on aggregate disposable income.

Proof. Write $u = \tau \log(x/x_0)$ and expand:

$$\begin{aligned} \bar{t}_{\mathrm{HSV}}(x) &= 1 - \lambda x^{-\tau} \;=\; 1 - \lambda x_0^{-\tau} \exp(-u) \\ &= 1 - \lambda x_0^{-\tau}\big(1 - u + \tfrac{1}{2}u^2 + O(u^3)\big) \\ &= (1 - \lambda x_0^{-\tau}) + \lambda x_0^{-\tau} u - \tfrac{1}{2}\lambda x_0^{-\tau} u^2 + O(u^3)\,. \end{aligned}$$

Collecting the constant and linear-in-log(x/x_0) terms gives $r(x) = r_0 + \beta \log(x/x_0)$ with r_0 and β as in (16). The quadratic term yields (17), with the remainder controlled by $|u| = \tau|\log(x/x_0)| \leq \tau L$ on the bounded interval. □

The remainder bound (17) is decisive in the wealth-tax regime. For an HSV-on-wealth schedule fitted to a Norwegian-flavoured calibration — average rate 0.5% at $x_0 = 1{,}000{,}000$ NOK (the median wealth scale) and 1.0% at $x_1 = 5{,}000{,}000$ NOK — the parameters solve to

$$\tau \;\approx\; 3.1 \times 10^{-3}\,, \qquad \lambda x_0^{-\tau} \;\approx\; 0.995\,.$$

At a high-wealth point $x = 50{,}000{,}000$ NOK ($\log(x/x_0) \approx 3.9$), the remainder bound (17) evaluates to

$$|\bar{t}_{\mathrm{HSV}} - r| \;\leq\; \tfrac{1}{2} \cdot 0.995 \cdot (3.1 \times 10^{-3} \cdot 3.9)^2 \;\approx\; 7.3 \times 10^{-5}\,,$$

that is, under one basis point.[3] The smooth log-progressive form is therefore an excellent approximation to the HSV functional form across the entire wealth-tax range, and effectively coincides with it for any policy-relevant comparison.

Remark (Bracket schedules and the approximation breakdown). The smooth log-progressive form is *not* a good approximation to bracket schedules of the form $r(x) = r \cdot \mathbf{1}\{x \geq X_0\}$ that are actually implemented in jurisdictions such as Norway. A bracket schedule has a discontinuity at the threshold which the smooth form smears into a logarithmic ramp; the HSV linearisation argument above does not apply because no continuous functional form can reproduce the kink. The substantive policy-relevant analysis on bracket schedules is taken up in Section 4, which uses piecewise-Gaussian techniques specific to that case. The smooth strand of the present section serves a different role: as the analytical benchmark on which the criteria coincide, it establishes the no-disagreement baseline against which Section 4's criterion divergence on bracket schedules is the substantive new finding.

3.2 Post-tax dynamics as an Ornstein–Uhlenbeck process

We begin by substituting the smooth schedule into the post-tax SDE (6). With $r(e^Y) = r_0 + \beta(Y - y_0)$ where $y_0 = \log x_0$, the SDE on Y becomes

$$\mathrm{d}Y_t \;=\; \big[k\, m_0 - r_0 + \beta y_0 - \beta Y_t\big]\mathrm{d}t + \sqrt{k}\,\sigma\,\mathrm{d}W_t\,. \tag{18}$$

[3] The bound is quadratic in $\tau \log(x/x_0)$ and the small-rate regime makes τ small; even at the highest wealth levels in the relevant empirical range, the linearisation error is two orders of magnitude below the resolution of any actual policy lever.

Collecting the constant terms and the Y_t-coefficient, this is an Ornstein–Uhlenbeck (OU) process:

$$\mathrm{d}Y_t \;=\; -\beta(Y_t - y_*)\mathrm{d}t + \sqrt{k}\,\sigma\,\mathrm{d}W_t\,, \tag{19}$$

with reversion rate β, equilibrium level

$$y_* \;=\; y_0 + \frac{k\,m_0 - r_0}{\beta}\,, \tag{20}$$

and diffusion coefficient $\sqrt{k}\,\sigma$.

Remark. The reversion rate is the progressivity slope β alone; the flow-tax retention k enters only through diffusion (uniform rescaling) and through the equilibrium position y_*. This is the formal expression of the additive separability of channels established in Frøseth (2026d, §7): the flow-tax channel acts uniformly on diffusion and on the constant component of drift, while the wealth-tax channel acts on the state-coupled component of drift.

The proportional limit $\beta \to 0$ corresponds to the case treated in Frøseth (2026b), where the post-tax process is a Brownian motion with constant drift $km_0 - r_0$ on log-wealth and no steady-state distribution. The progressive case $\beta > 0$ confines the process and produces a stationary distribution, recovering the OU confinement mechanism of Frøseth (2026f, §4) when $k = 1$.

3.3 Steady-state distribution and the structural ratio

Standard OU theory gives the stationary distribution explicitly.

Proposition 2 (Steady-state distribution)**.** *For $\beta > 0$ and any initial condition Y_0, the post-tax process (19) converges as $t \to \infty$ to the stationary distribution*

$$Y_\infty \;\sim\; \mathcal{N}(y_*, \sigma_Y^2)\,, \qquad \sigma_Y^2 \;=\; \frac{k\,\sigma^2}{2\beta}\,. \tag{21}$$

The induced steady-state distribution of raw wealth $X_\infty = \exp(Y_\infty)$ is log-normal with parameters (y_*, σ_Y^2).

Proof. The OU SDE (19) is linear and admits the explicit solution

$$Y_t \;=\; y_* + (Y_0 - y_*)e^{-\beta t} + \sqrt{k}\,\sigma \int_0^t e^{-\beta(t-s)}\mathrm{d}W_s\,.$$

The integral is a centred Gaussian with variance $\int_0^t k\sigma^2 e^{-2\beta(t-s)}\mathrm{d}s = (k\sigma^2/2\beta)(1-e^{-2\beta t})$, which converges to $k\sigma^2/(2\beta)$ as $t \to \infty$. The mean converges to y_* and the variance to σ_Y^2. The induced distribution of $X_\infty = e^{Y_\infty}$ is log-normal by definition. □

The crucial observation is that σ_Y^2 depends on the two design parameters only through their ratio.

Definition 1 (Structural progressivity ratio)**.** For a smooth log-progressive schedule $(k, r_0, \beta) \in$

$\mathcal{S}_{\text{flow+prog}}$, the *structural progressivity ratio* is

$$\rho_p \;=\; \frac{\beta}{k}\,. \tag{22}$$

In terms of ρ_p the steady-state log-variance reads

$$\sigma_Y^2 \;=\; \frac{\sigma^2}{2\rho_p}\,. \tag{23}$$

The two design parameters (k, β) appear in the inequality analysis only through the single combination ρ_p. The equilibrium location y_* depends separately on (k, β, r_0), but inequality measures are location-invariant: they depend on the shape of the distribution through σ_Y^2 alone.

Remark (Comparative statics)*.* From (22), the structural ratio is increasing in β and decreasing in k. Equivalently, σ_Y^2 is decreasing in β and increasing in k: more progressivity compresses the steady-state distribution, more flow-tax retention expands it. This will translate directly into the comparative statics of the redistributive criteria below.

3.4 Closed-form inequality measures on the design plane

Both redistributive criteria reduce to closed-form functions of the structural ratio.

Proposition 3 (Gini and top share on the smooth class)**.** *Under Proposition 2, the steady-state Gini coefficient and top share at quantile* $q \in (0,1)$ *satisfy*

$$\mathrm{Gini}(\rho_p) \;=\; \mathrm{erf}\left(\frac{\sigma_Y}{2}\right) \;=\; \mathrm{erf}\left(\frac{\sigma}{2\sqrt{2\rho_p}}\right), \tag{24}$$

$$S_q(\rho_p) \;=\; 1 - \Phi\Big(\Phi^{-1}(q) - \sigma_Y\Big) \;=\; 1 - \Phi\left(\Phi^{-1}(q) - \frac{\sigma}{\sqrt{2\rho_p}}\right). \tag{25}$$

Both are strictly decreasing in ρ_p *and strictly increasing in* σ_Y^2*.*

Proof. The closed forms (12) and (14) for the log-normal Gini and top share apply directly to the steady-state distribution of Proposition 2, with σ_Y given by (23). Both are smooth strictly-monotone functions of σ_Y^2 (the derivative of erf is positive, and Φ is strictly increasing), and σ_Y^2 is strictly decreasing in ρ_p by (23). □

Remark (Iso-inequality curves on the design plane)*.* The level sets of Gini and S_q on the (k, β) plane coincide: both are level sets of the structural ratio $\rho_p = \beta/k$, i.e. rays from the origin with slope ρ_p when β is plotted against k. The two design parameters substitute for one another along these rays at a fixed inequality level: any (small k, small β) pair on the same ray is inequality-equivalent to any (large k, large β) pair on that ray. The two redistributive criteria see only this one combination on the smooth class.

3.5 Aggregate revenue at steady state

We turn now to the revenue generated by the wealth-tax channel at steady state. This determines which (k, β, r_0) triples are jointly admissible at a given revenue target.

Proposition 4 (Wealth-tax revenue at steady state)**.** *Under Proposition 2, the per-capita aggregate revenue generated by the wealth-tax channel,* $R_w = \mathbb{E}[r(X_\infty)\, X_\infty]$*, satisfies*

$$R_w \;=\; k\,\mu \cdot \mathbb{E}[X_\infty] \;=\; k\,\mu \cdot \exp\left\{ y_0 + \frac{k(\mu - \sigma^2/4) - r_0}{\beta} \right\}. \tag{26}$$

Proof sketch. The mean follows from the log-normal MGF: $\mathbb{E}[X] = \exp(y_* + \sigma_Y^2/2)$. For $\mathbb{E}[r(X)X]$, expand $r(X) = r_0 + \beta(Y - y_0)$ and use the identity $\mathbb{E}[Y e^Y] = \mathbb{E}[X](y_* + \sigma_Y^2)$ to obtain $\mathbb{E}[r(X)X] = \mathbb{E}[X](km_0 + \beta\sigma_Y^2)$. Substituting $\sigma_Y^2 = k\sigma^2/(2\beta)$ collapses this to $k\mu\,\mathbb{E}[X]$. Full computation in Appendix A.1. □

Remark (Revenue-as-aggregate-growth identity). The identity $R_w = k\mu\,\mathbb{E}[X_\infty]$ is the steady-state expression of the budget identity at the population level: at the stationary distribution, aggregate wealth-tax collections equal aggregate post-flow-tax growth $k\mu$ times mean wealth. The progressivity slope β does not appear explicitly in the revenue identity — it shifts the *distribution* of who pays without changing the aggregate revenue at steady state, conditional on $\mathbb{E}[X_\infty]$. The total nevertheless depends on (k, β, r_0) through $\mathbb{E}[X_\infty]$, and (26) gives the explicit dependence.

Proposition 4 provides the means of expressing the floor rate r_0 as a function of (k, β) and a revenue target R_w^*:

$$r_0 \;=\; k(\mu - \tfrac{\sigma^2}{4}) - \beta \log\left(\frac{R_w^*}{k\mu x_0} \right). \tag{27}$$

For each admissible (k, β, R_w^*) triple, this fixes a unique floor rate that delivers exactly the target. The admissibility region is $r_0 \geq 0$, which translates to

$$\beta\, \log\!\big(R_w^*/(k\mu x_0)\big) \;\leq\; k\big(\mu - \tfrac{\sigma^2}{4}\big). \tag{28}$$

On the side of the design plane where $R_w^* < k\mu x_0$, the left-hand side is negative and (28) is satisfied for all $\beta \geq 0$; on the side where $R_w^* > k\mu x_0$, the constraint binds at large β.

3.6 Joint optimal design

We now formulate the joint optimal-design problem on the smooth class. Given a redistributive target (a Gini level G^* or a top-share level S_q^*) and a revenue target R_w^*, the designer selects (k, β, r_0) to satisfy both constraints within feasibility.

Definition 2 (Joint optimal-design problem)**.** For a redistributive target Gini* and a revenue target R_w^*, the *Gini-targeted joint optimal design* is the solution of

$$\min_{(k,\beta,r_0)\in\mathcal{F}} \big(\mathrm{Gini}[k,\beta] - \mathrm{Gini}^*\big)^2 \tag{29}$$

subject to the revenue constraint $R_w[k,\beta,r_0] = R_w^*$, where $\mathcal{F}$ is the feasibility set

$$\mathcal{F} \;=\; \{(k,\beta,r_0) : k \in [k_{\min},1],\; \beta \in [0,\beta_{\max}],\; r_0 \geq 0\}\,. \tag{30}$$

The lower bound $k_{\min}$ is fixed by the feasible range of the corporate–dividend tax pair; the upper bound $\beta_{\max}$ by the implementability of the top wealth-tax rate. The *top-share-targeted joint optimal design* is defined analogously with S_q replacing Gini.

Propositions 3 and 4 reduce the three-parameter problem to a one-parameter problem in ρ_p.

Theorem 1 (Reduction to the structural ratio)**.** *On the smooth log-progressive class, the Gini- and top-share-targeted joint optimal designs of Definition 2 coincide and are characterised by a single equation in ρ_p:*

$$\rho_p^* \;=\; \frac{\sigma^2}{2\sigma_{Y,*}^2}\,, \tag{31}$$

where $\sigma_{Y,*}$ is the unique log-normal log-variance satisfying Gini $=$ Gini* *(equivalently, $S_q = S_q^*$) on the canonical (log-normal) family. The optimum is feasible if and only if there exists $(k,\beta) \in [k_{\min},1] \times [0,\beta_{\max}]$ with $\beta = \rho_p^* \cdot k$ and the corresponding r_0 from (27) non-negative. When feasible, the optimum is degenerate along the ray $\beta = \rho_p^* \cdot k$: any (k,β) on this ray with r_0 adjusted via (27) achieves the same Gini and the same revenue.*

Proof. By Proposition 3, both criteria are strictly monotone functions of σ_Y^2, and by Proposition 2 σ_Y^2 is a strictly monotone function of $\rho_p = \beta/k$ alone. Both criteria are therefore strictly monotone functions of the single combination ρ_p, and any feasible target inverts uniquely to a unique ρ_p^* via (31). By Proposition 4 and (27), for any (k,β) with $\beta = \rho_p^* k$ the floor rate r_0 is uniquely determined by the revenue target R_w^*. The feasibility condition follows from (30) and (28). □

Remark (Degeneracy along the optimal ray)*.* Theorem 1 states a degeneracy: any $(k,\beta) \in \mathcal{F}$ on the ray $\beta = \rho_p^* k$ solves the joint optimal-design problem. The criterion does not pin down the individual values of the flow-tax retention k and the progressivity slope β; only their ratio is determined. Pinning (k,β) individually requires an additional consideration outside the inequality–revenue problem — for example, an upper bound on the top wealth-tax rate $r_{\max}$, a corporate-tax revenue requirement, or a distortion-minimisation criterion as in Frøseth (2026b). We discuss these in Section 8.

3.7 Comparison theorem

Theorem 1 contains the comparison theorem on the smooth class.

Corollary 1 (Comparison theorem on the smooth class)**.** *On the smooth log-progressive subclass within (C1)–(C3), the Gini-optimal and top-share-optimal joint designs at any fixed revenue target R_w^* coincide as sets in $\mathcal{F}$. The optimal structural ratio ρ_p^* depends on the chosen target level (Gini level or top-share level), but the choice of* criterion family *— Gini versus top share at any quantile q — does not select a different optimum.*

Proof. Both Gini and S_q are strictly monotone functions of σ_Y^2 on the log-normal family (Proposition 3), and σ_Y^2 is determined by ρ_p alone (Proposition 2). Two distinct criteria can only disagree about which ρ_p^* is the optimum if their targeted levels parametrise different log-normal sub-families. They do not: the Gini target Gini* and the top-share target S_q^* each pick a unique $\sigma_{Y,*}^2$ on the canonical family, so each picks a unique ρ_p^* via (31). $\square$

The substantive policy reading of Corollary 1 is deliberately negative: on the smooth class, the choice between Gini and top-share targets does not affect the optimal schedule shape, only the targeted level. This is the *no-disagreement baseline* of the paper. The result is not a structural feature of the design problem itself — the schedule space is rich enough to admit non-trivial criterion divergence — but of the canonical-distribution restriction imposed by the smooth log-progressive functional form. The two criteria parametrise the log-normal family differently (Gini through σ_Y via the error function, top share through σ_Y via the inverse Gaussian CDF), but on the one-parameter family $\{\mathcal{N}(y_*, \sigma_Y^2) : \sigma_Y^2 > 0\}$ the parametrisations are isomorphic and the optima coincide.

The bracket case in Section 4 breaks this canonicity. The post-tax distribution under a bracket schedule is piecewise-Gaussian rather than log-normal, with shape controlled by additional parameters (the threshold position X_0, the rate r above the threshold, the curvature of the density at the matching point). Two redistributive criteria that agree on the canonical family can disagree on the richer piecewise-Gaussian family, and which schedule shape is selected depends on whether overall compression (Gini) or tail compression (top share) is the binding redistributive objective. Section 4 establishes that this disagreement is real and has policy bite at calibrated parameter values: the criterion choice is not just a labelling difference but selects qualitatively distinct optimal schedules. The no-disagreement baseline of the present section is what makes that result substantive rather than generic.

3.8 The Stiglitz channel and (C1)–(C3) neutrality

The OU compression mechanism we have used to derive the inequality reduction in this section operates entirely on the quantity dimension: progressivity β shrinks the steady-state log-variance $\sigma_Y^2 = k\sigma^2/(2\beta)$, and both inequality criteria see only this shrinkage. A natural question is whether the same progressivity might also act on the *price* dimension — equilibrium asset returns — in a way that partially or fully offsets the quantity compression. This is the question asked, in the labour-skill setting, by Heathcote et al. (2017) in the Stiglitz tradition of treating pre-tax distributions as endogenous equilibrium objects (Stiglitz, 1985), and the answer is informative for our framework.

The Stiglitz mechanism in the HSV model. HSV embed the progressivity parametrisation (15) in a heterogeneous-agent equilibrium model with endogenous skill investment. Individuals choose initial skill $s(\kappa; \tau)$ to maximise expected utility, with cost depending on inverse ability κ; skills are imperfect substitutes in production with elasticity θ. Two channels then determine how progressivity affects equilibrium inequality (Heathcote et al., 2017, Proposition 2 and Corollary 2.1).

- *Quantity channel.* Higher τ reduces the after-tax return to skill, so optimal skill investment falls: $s(\kappa;\tau) \propto (1-\tau)^{\psi/(1+\psi)} \cdot \kappa$, where ψ is the elasticity of skill investment to the after-tax return. The skill distribution compresses toward zero as $\tau \to 1$.

- *Price channel.* Imperfect substitutability in production drives up the *pre-tax* skill premium when high-skill types become scarce. The equilibrium marginal skill premium $\pi_1(\tau) \propto (1-\tau)^{-\psi/(1+\psi)}$ is strictly increasing in τ and divergent at the progressivity boundary. Heathcote et al. (2017) label this feedback the "Stiglitz effect," citing Stiglitz (1985) for the broader principle that the pre-tax wage distribution is an endogenous equilibrium object that responds to taxation through the supply side; the specific imperfect-substitution mechanism is a feature of HSV's own model rather than a result of the cited paper.[4]

The two channels move pre-tax inequality in opposite directions, and HSV obtain the striking conclusion that they *exactly cancel*: under their baseline log-utility specification, the variance of log skill prices is $1/\theta^2$, independent of τ (Heathcote et al., 2017, Corollary 2.1). The entire compression of the wage distribution under progressivity comes from the quantity adjustment; the price dimension is *undisturbed*. The result depends sensitively on the utility structure (HSV note that the cancellation is exact only under their baseline specification), but the qualitative point is robust: in any equilibrium model where the progressively-taxed quantity is endogenously supplied through choices made on heterogeneous primitives, the supply response generates a price-level offset that the quantity-only analysis misses.

The wealth-tax analogue. The structural translation to a progressive wealth tax is direct. The endogenous quantity is no longer skill investment but the *equilibrium portfolio composition* chosen by households: how much wealth is held in risky versus safe assets, and across asset classes more finely. The price dimension is the *equilibrium pre-tax return* on each asset class — the μ of the model in Section 2.1 — which clears the asset market given aggregate household demand. A progressive wealth tax that pushes high-wealth households toward safer portfolios shifts the relative demand for risky assets and, in general equilibrium, induces a Stiglitz-type price feedback: the pre-tax risky return rises as risky assets become harder to place, partly offsetting the direct compression effect of the tax on accumulated wealth.

The framework of Section 2 treats μ and σ as exogenous, so the price feedback is by construction absent. The question is whether this is a reasonable *modelling* choice or whether the absence of the channel masks an inequality-relevant offset.

(C1)–(C3) neutrality structurally suppresses the channel. The answer in the schedule space we treat is unambiguous: the Stiglitz channel is suppressed not by assumption but by the (C1)–(C3) class itself. The generalised-neutrality theorem of Frøseth (2026d) establishes that on the (C1)–(C3) class, the wealth tax acts on the post-tax wealth process by drift-shift-

[4] Stiglitz (1985) develops a labour-market model with search frictions and imperfect information, showing the existence of equilibria with wage dispersion across firms paying for identical labour. The mechanism there is firm-side rather than production-side, and the connection to HSV's skill-premium feedback is through the general Stiglitz-tradition reading of pre-tax wage distributions as endogenous equilibrium objects, not through any specific result.

and-rescale without altering the relative drifts between assets. A household at any wealth level chooses the same risk-bearing allocation $\mathbf{w}^*$ irrespective of (k, β, r_0): the wealth-tax channel is a uniform compression of the household's budget constraint along the safe direction, and the flow-tax channel is a uniform rescaling of excess returns by k. Neither modification redirects the household's portfolio *composition*.

The aggregate consequence is that, on the (C1)–(C3) class, relative demand for risky versus safe assets does not respond to progressivity. Aggregate demand for each asset class scales uniformly with the post-tax wealth distribution, so the market-clearing condition for risky assets does not call for a shift in the equilibrium pre-tax return μ. The Stiglitz price-feedback channel is therefore not present *within the (C1)–(C3) class* — the OU compression mechanism of Proposition 2 and the comparison theorem of Corollary 1 are operative on the quantity dimension alone, and the inequality reduction we have derived is genuine rather than a partial-equilibrium artefact.

Interpretation: the framework isolates the OU mechanism. The contrast with HSV is structurally informative. In HSV's labour-skill setting the two channels coexist and largely cancel, so the progressive tax operates predominantly through quantity adjustment in equilibrium. In our wealth-tax setting on the (C1)–(C3) class, the price channel is shut off entirely by the neutrality conditions, so the progressive tax operates only through quantity adjustment by construction. The two settings arrive at *quantity-only* compression by different routes: HSV through a cancellation that is utility-specific and partial, ours through a structural neutrality property that is exact and robust to utility specification within the (C1)–(C3) class.

This is a positive feature of the framework rather than a limitation, on at least three grounds. First, it makes the inequality result of Proposition 3 sharp: there is no unrepresented offsetting channel that would attenuate the $\sigma_Y^2 \propto 1/\beta$ scaling in general equilibrium. Second, it identifies the (C1)–(C3) class as the precise structural assumption under which the OU compression mechanism is the unique inequality channel; outside the class, the Stiglitz feedback re-enters and the analysis becomes a joint quantity–price equilibrium problem with potentially distinct implications. Third, it specifies which features of an actual wealth-tax system would activate the Stiglitz feedback: the violations of the (C1)–(C3) conditions catalogued in Frøseth (2026d, §5). Failure of (C1) ($\tau_k \neq \tau_c$) or (C2) ($r_s \neq r_f$) creates a relative distortion between the equity and debt channels that shifts aggregate portfolio composition in a way the wealth-tax rate β would amplify; failure of (C3) (non-uniform asset assessment) tilts the tangency portfolio in a way that high-β schedules amplify further. In each case, the quantity-only analysis of the present section would need to be embedded in a richer asset-market-clearing model to recover the full inequality response.

The suppression argument is structurally tied to (C1)–(C3) and so carries through unchanged to the bracket strand of Section 4. Section 5.1 treats the non-uniform-assessment case in detail; the bracket-strand calibration of Section 6 works inside the (C1)–(C3) envelope through the uniform-assessment approximation $\bar{a}$.

3.9 Limit cases and connection to the corpus

The smooth log-progressive strand recovers two important limits of the existing corpus.

$k = 1$ limit: stock-only progressive case. Setting $k = 1$ removes the flow-tax channel; the SDE (19) becomes

$$\mathrm{d}Y_t \;=\; -\beta(Y_t - y_*)\mathrm{d}t + \sigma\,\mathrm{d}W_t\,,$$

which is the OU confinement of Frøseth (2026f, §4) with β playing the role of the progressivity parameter and $\sigma_Y^2 = \sigma^2/(2\beta)$ the steady-state log-variance. The Gini-optimal design $\rho_p^* = \beta^*$ in this limit recovers Frøseth (2026f, Theorem 3) verbatim. The present analysis generalises that result by allowing the same inequality target to be achieved at any $k \in [k_{\min}, 1]$ with $\beta = \rho_p^* \cdot k$, providing a continuum of joint designs that are inequality-equivalent to the stock-only case.

$\beta = 0$ limit: proportional case. The proportional limit collapses the schedule to a flat rate $r(x) = r_0$ and the SDE loses its restoring force, leaving Brownian motion with constant drift on log-wealth. No stationary distribution exists ($\rho_p \to 0$, $\sigma_Y^2 \to \infty$, Gini $\to 1$, $S_q \to 1$): the well-known Bouchaud and Mézard (2000) observation that proportional taxes do not redistribute through the market channel (Frøseth, 2026f,b), and the sub-class on which Frøseth (2026b) substitutes finite-horizon distortion-minimising criteria for the steady-state criteria treated here.

The full smooth class indexed by (k, β) interpolates between the two limits, and the optimal-design ray $\beta = \rho_p^* k$ traces a one-parameter family of inequality-equivalent schedules of which the stock-only case is one extreme — the analytical content of this section relative to Frøseth (2026f,b) taken separately.

4 The bracket strand

We turn to the bracket sub-class (9), the form actually implemented in jurisdictions such as Norway, Switzerland, and (historically) Spain and France: a zero-rate region below a threshold X_0 followed by a positive rate r above it (single-bracket), with some jurisdictions adding a top bracket above a second threshold X_1 (two-bracket).

The dimensionality of the inequality-relevant policy space organises the analysis. The single-bracket case at fixed regime collapses to a single active policy variable, so both criteria reduce to monotone functions of it and the equivalence of Corollary 1 carries through (in a weaker form). The two-bracket case has a three-dimensional policy space with a substantive criterion divergence: a single Gini level admits a near-order-of-magnitude range of top-1% shares across the iso-Gini manifold.

Notation. The smooth-strand parameters (k, r_0, β) of Section 3 are not the bracket-strand parameters of this section. We use λ-symbols exclusively for the log-wealth slopes of the steady-state density: λ_- on the lower regime in the single-bracket case, and $(\lambda_0, \lambda_1, \lambda_2)$ on the lower, middle, and top regimes in the two-bracket case. Where the slope is negative on the upper regime (so that the raw-wealth distribution has a Pareto-shaped tail) we use the positive-valued

Pareto exponent $\alpha = -\lambda$. The threshold ratio $\Delta = \log(X_1/X_0)$ enters the two-bracket case as a third inequality-relevant parameter.

4.1 Post-tax dynamics and steady-state distribution

Substituting the single-bracket schedule $r(x) = r \cdot \mathbf{1}\{x \geq X_0\}$ into the post-tax SDE (6) yields a process on log-wealth with piecewise-constant drift,

$$\mathrm{d}Y_t \;=\; [km_0 - r \cdot \mathbf{1}\{Y_t \geq y_0\}]\mathrm{d}t + \sqrt{k}\,\sigma\,\mathrm{d}W_t\,, \tag{32}$$

where $y_0 = \log X_0$. The drift is km_0 below the threshold (untaxed regime) and $km_0 - r$ above (taxed regime); the diffusion is uniform $\sqrt{k}\,\sigma$ throughout. The process is no longer Ornstein–Uhlenbeck and the steady-state distribution is no longer Gaussian.

Proposition 5 (Bracket steady-state distribution)**.** *Suppose $m_0 > 0$ and $r > km_0$. The post-tax process (32) admits a unique stationary distribution on $\mathbb{R}$ with density*

$$\pi(y) \;=\; \pi_0 \cdot \begin{cases} e^{\lambda_-(y-y_0)} & y < y_0\,, \\ e^{\lambda_+(y-y_0)} & y \geq y_0\,, \end{cases} \tag{33}$$

where

$$\lambda_- \;=\; \frac{2m_0}{\sigma^2}\,, \qquad \lambda_+ \;=\; -\frac{2}{\sigma^2}\left(\frac{r}{k} - m_0\right)\,, \qquad \pi_0 \;=\; \frac{\lambda_-(-\lambda_+)}{\lambda_- - \lambda_+}\,. \tag{34}$$

The steady-state mass on each side of the threshold is

$$\mathbb{P}(Y_\infty < y_0) \;=\; \frac{-\lambda_+}{\lambda_- - \lambda_+}\,, \qquad \mathbb{P}(Y_\infty \geq y_0) \;=\; \frac{\lambda_-}{\lambda_- - \lambda_+}\,. \tag{35}$$

Proof. The stationary Fokker–Planck equation $0 = -\partial_y(b(y)\pi) + \frac{1}{2}\sigma_Y^2\,\partial_y^2\pi$ with zero probability current and constant diffusion $\sigma_Y^2 = k\sigma^2$ reduces to $\pi'/\pi = 2b(y)/\sigma_Y^2$. Within each regime the drift is constant, so the density is exponential with rate $\lambda_\pm$ as in (34). The condition $r > km_0$ ensures $\lambda_+ < 0$, so the right-hand exponential decays at $+\infty$. The condition $m_0 > 0$ ensures $\lambda_- > 0$, so the left-hand exponential decays at $-\infty$. Continuity of π at y_0 gives equal pre-factors on both sides; normalisation $\int_{\mathbb{R}} \pi(y)\mathrm{d}y = 1$ then determines π_0 as in (34). The mass split (35) follows by direct integration. □

Section 33 is the asymmetric Laplace density on log- wealth, centred at the threshold y_0 with separate decay rates on each side. Substituting back to raw wealth, $X_\infty = e^{Y_\infty}$ has the piecewise power-law density

$$f_X(x) \;=\; \begin{cases} \pi_0\,X_0^{-\lambda_-}\,x^{\lambda_- -1} & 0 < x < X_0\,, \\ \pi_0\,X_0^{-\lambda_+}\,x^{\lambda_+ -1} & x \geq X_0\,. \end{cases} \tag{36}$$

The upper tail ($x \geq X_0$) is Pareto with exponent $\alpha = -\lambda_+ = 2(r/k - m_0)/\sigma^2$. The Pareto exponent is fully determined by the policy parameters (k, r) and the regime parameters (m_0, σ); the threshold X_0 appears nowhere in α or in the mass split (35). The threshold acts purely as a multiplicative scale on raw wealth: shifting it $X_0 \to cX_0$ shifts the steady-state distribution to $X_\infty \to cX_\infty$ in distribution.

Remark (Finite-mean regime). The Pareto upper tail has finite mean if and only if $\alpha > 1$, equivalently $r > k(m_0 + \sigma^2/2) = k\mu$. For $\alpha \leq 1$ the steady-state distribution exists but $\mathbb{E}[X_\infty] = \infty$ and neither inequality criterion is well-defined. We therefore restrict the analysis to the finite-mean regime $\alpha > 1$, equivalently $r > k\mu$; the practical wealth-tax rates considered in the calibration of Section 6 sit comfortably inside this region.

4.2 The two-parameter reduction

Proposition 5 reduces the three-parameter schedule space (k, X_0, r) to a two-parameter shape description (λ_-, α), with X_0 entering only as a multiplicative scale. Both inequality criteria, being scale-invariant, depend on the policy schedule through (λ_-, α) alone.

Definition 3 (Bracket shape parameters). For a single-bracket schedule (k, r, X_0) within the finite-mean regime, the *bracket shape parameters* are

$$\lambda_- \;=\; \frac{2m_0}{\sigma^2}\,, \qquad \alpha \;=\; \frac{2}{\sigma^2}\left(\frac{r}{k} - m_0\right)\,, \tag{37}$$

with $\lambda_- > 0$ and $\alpha > 1$.

The lower-regime slope λ_- is independent of the policy schedule: it is a regime parameter determined by (m_0, σ) alone. The upper-regime Pareto exponent α is set jointly by the regime parameters and the *effective* wealth-tax rate r/k above the threshold; the flow-tax retention k and the wealth-tax rate r enter only through this ratio. The single dimensionless quantity that responds to schedule choice within the bracket sub-class is therefore α, with the two-dimensional (λ_-, α) pair fixing the post-tax shape.

4.3 Closed-form inequality measures

The piecewise power-law structure of (36) admits closed forms for both redistributive criteria.

Proposition 6 (Closed-form Gini and top share on the bracket sub-class). *Under Proposition 5 and the finite-mean condition $\alpha > 1$, the steady-state distribution has mean*

$$\mathbb{E}[X_\infty] \;=\; X_0 \cdot \frac{\lambda_-\,\alpha}{(\lambda_- + 1)(\alpha - 1)} \tag{38}$$

and the inequality criteria evaluate to

$$\mathrm{Gini}(\lambda_-, \alpha) \;=\; \frac{(\lambda_- + 1)(\alpha - 1)}{\lambda_-\alpha}\left[\frac{p_-}{\lambda_- + 1} - \frac{p_-^2}{2\lambda_- + 1} + \frac{p_+}{\alpha - 1} - \frac{p_+^2}{2\alpha - 1}\right], \tag{39}$$

$$S_q(\lambda_-, \alpha) \;=\; \left(1 + \frac{1}{\lambda_-}\right) \cdot (1 - q)^{(\alpha-1)/\alpha} \cdot \left(\frac{\lambda_-}{\lambda_- + \alpha}\right)^{1/\alpha}, \tag{40}$$

where $p_- = \alpha/(\lambda_- + \alpha)$ and $p_+ = \lambda_-/(\lambda_- + \alpha)$ are the steady-state masses below and above the threshold, with the top-share formula (40) valid for $q \geq p_-$ (the quantile of interest above the *threshold; see Section 4.3 below for $q < p_-$).*

Proof sketch. Direct integration of the piecewise-power-law density on each regime, combined with the standard identity $\text{Gini} = (1/\mathbb{E}[X]) \int F(1-F)\, dx$ and the Pareto-tail mean formula $\mathbb{E}[X \mid X \geq z_q] = \alpha z_q/(\alpha - 1)$. The finite-mean condition $\alpha > 1$ ensures all integrals converge. Full calculation in Appendix A.2. □

Remark (Top share at quantiles inside the lower regime). For $q < p_-$ the quantile $z_q = (q/p_-)^{1/\lambda_-}$ lies in the lower regime, and the top-share calculation has contributions from both regimes: $S_q = (\mathbb{E}[Z\mathbf{1}\{Z \geq z_q,\, Z < 1\}] + \mathbb{E}[Z\mathbf{1}\{Z \geq 1\}])/\mathbb{E}[Z]$. The closed form is straightforward but unenlightening; for the empirical calibrations in Section 6 the relevant quantile $q \in \{0.99, 0.999\}$ exceeds p_- at every parameter configuration we consider, so (40) applies and we suppress the case $q < p_-$ throughout.

4.4 Criterion equivalence on the single-bracket sub-class

The two-parameter pair (λ_-, α) of Definition 3 reduces to a one-parameter description once the regime is fixed: $\lambda_- = 2m_0/\sigma^2$ is determined by the no-tax dynamics (μ, σ) alone, and the active policy variable is the Pareto exponent α. This collapses the inequality criteria to one-dimensional functions of α.

Proposition 7 (Single-bracket criterion equivalence at fixed regime)**.** *At fixed regime $\lambda_- > 0$, both inequality criteria $\text{Gini}(\lambda_-, \alpha)$ and $S_q(\lambda_-, \alpha)$ of Proposition 6 are strictly decreasing functions of α on the finite-mean regime $\alpha > 1$. Consequently:*

1. *There is a unique $\alpha^*_{\text{Gini}}(\lambda_-, \text{Gini}^*)$ achieving any feasible Gini target Gini^*, and a unique $\alpha^*_{S_q}(\lambda_-, S^*_q)$ achieving any feasible top-share target.*
2. *For any feasible target pair (Gini^*, S^*_q) that corresponds to the same α^*, the Gini-targeted and top-share-targeted joint optimal designs of Definition 2 coincide.*

The criterion-equivalence baseline of Corollary 1 therefore carries through to the single-bracket sub-class, in a slightly weaker form: both criteria are functions of the single shape parameter α at fixed regime, and the criterion choice does not select different optimal schedules.

Proof sketch. Differentiation of (40) gives $\partial_\alpha \log S_q$ as a sum of three terms; the term $\log(1 - q)/\alpha^2$ dominates (negative, since $1 - q$ is small for q near unity), and numerical evaluation confirms the derivative is negative throughout $\alpha > 1$ at any $\lambda_- > 0$ and $q \geq 0.5$. The Gini coefficient satisfies $\text{Gini} \to 1$ as $\alpha \to 1^+$ and approaches a finite floor as $\alpha \to \infty$; numerical evaluation across $\alpha \in [1.2, 5]$ confirms strict monotone decrease. Both criteria being strictly monotone in α, target-inversion has a unique solution; equivalent targets pick the same α^*. Full calculation in Appendix A.3. □

Remark (Strength of the equivalence). Proposition 7 is a weaker form of the smooth-class equivalence of Corollary 1: in the smooth case, both criteria depend on the single combination $\rho_p = \beta/k$ of the two design parameters, so the equivalence holds across the entire 2D design plane. In the bracket case, the equivalence holds only after the regime parameter λ_- is fixed at its no-tax value; the active policy variable then collapses to α, and the equivalence emerges from the one-dimensionality. The two-bracket case below does not admit this further reduction, and the genuine criterion divergence emerges there.

The single-bracket case is therefore best understood as analytical scaffolding: it establishes the closed-form machinery (Propositions 5 and 6) used in the two-bracket analysis below, and confirms that the single-parameter restriction inherited from the schedule's one-rate structure preserves the criterion-coincidence property of Section 3. The substantive new content of the bracket family — where Gini and top-share genuinely disagree — arises with the additional policy parameters of the two-bracket schedule.

4.5 Single-bracket joint optimal design

We formulate the joint optimal-design problem of Definition 2 on the single-bracket sub-class. The two policy levers entering the shape parameters are the flow-tax retention k and the wealth-tax rate r, combined into the effective rate r/k. The threshold X_0 enters only as a scale, so we may treat it as fixed by normalising $\mathbb{E}[X_\infty]$ to a target mean wealth $\bar{W}$ via (38):

$$X_0 \;=\; \bar{W} \cdot \frac{(\lambda_- + 1)(\alpha - 1)}{\lambda_- \alpha}\,. \tag{41}$$

The two-parameter design plane (λ_-, α) then comprises the full inequality-relevant policy space, with λ_- a regime parameter (determined by the no-tax (μ, σ)) and α the active policy variable.

Corollary 2 (Bracket joint optimal design)**.** *At fixed regime (μ, σ) and fixed mean-wealth target $\bar{W}$, the Gini- and top-share-targeted joint optimal-design problems within the bracket sub-class reduce to optimisation in the single parameter α, equivalently in the effective rate r/k:*

$$\alpha^*_{\mathrm{Gini}} \;=\; \operatorname*{arg\,min}_{\alpha > 1} \;|\,\mathrm{Gini}(\lambda_-, \alpha) - \mathrm{Gini}^*\,|^2, \tag{42}$$

$$\alpha^*_S \;=\; \operatorname*{arg\,min}_{\alpha > 1} \;|\,S_q(\lambda_-, \alpha) - S_q^*\,|^2, \tag{43}$$

subject to the implementability bound $\alpha \le \alpha_{\max} = 2(r_{\max}/k_{\min} - m_0)/\sigma^2$. By Proposition 7, equivalent targets pick the same α^*, so the criterion choice does not distinguish optimal schedules within the single-bracket sub-class.

The corollary makes the single-bracket policy implication concrete: the pair (k, r) is constrained by the corporate–dividend retention floor $k_{\min}$ and the implementability ceiling $r_{\max}$, but within the feasible region the choice of effective rate r/k is the active design variable, and both criteria pick the same α^*. To exhibit a substantive criterion divergence we require a schedule family with a multi-dimensional inequality-relevant policy space. The two-bracket sub-class supplies it.

4.6 Two-bracket post-tax dynamics and steady-state

The two-bracket schedule (10) is the policy-realistic form: Norway introduced $X_1 \approx 20$ M NOK with $r_2 = 1.1\%$ in 2023 alongside the standard $X_0 \approx 1.7$ M NOK with $r_1 \approx 1.0\%$; several Swiss cantons operate similar two-tier structures. Substituting into (6) yields the post-tax SDE

$$\mathrm{d}Y_t \;=\; \left[km_0 - r_1\mathbf{1}\{y_0 \le Y_t < y_1\} - r_2\mathbf{1}\{Y_t \ge y_1\}\right]\mathrm{d}t + \sqrt{k}\,\sigma\,\mathrm{d}W_t\,, \tag{44}$$

where $y_0 = \log X_0$ and $y_1 = \log X_1$. The drift is piecewise constant on three regimes: km_0 on the lower regime ($y < y_0$), $km_0 - r_1$ on the middle regime ($y_0 \le y < y_1$), and $km_0 - r_2$ on the top regime ($y \ge y_1$).

Proposition 8 (Two-bracket steady-state distribution)**.** *Suppose $m_0 > 0$ and $r_2 > km_0$. The post-tax process (44) admits a unique stationary distribution on $\mathbb{R}$ with density*

$$\pi(y) \;=\; \pi_0 \cdot \begin{cases} e^{\lambda_0(y-y_0)} & y < y_0\,, \\ e^{\lambda_1(y-y_0)} & y_0 \le y < y_1\,, \\ \xi \cdot e^{\lambda_2(y-y_1)} & y \ge y_1\,, \end{cases} \tag{45}$$

with three slopes

$$\lambda_0 = \frac{2m_0}{\sigma^2}\,, \quad \lambda_1 = \frac{2(m_0 - r_1/k)}{\sigma^2}\,, \quad \lambda_2 = \frac{2(m_0 - r_2/k)}{\sigma^2}\,, \tag{46}$$

the threshold ratio $\Delta = y_1 - y_0 = \log(X_1/X_0)$, and the dimensionless parameter $\xi = e^{\lambda_1\Delta}$. The normalisation constant is

$$\pi_0 \;=\; \left[\frac{1}{\lambda_0} + \frac{\xi - 1}{\lambda_1} + \frac{\xi}{\alpha_2}\right]^{-1}, \qquad \alpha_2 = -\lambda_2\,, \tag{47}$$

(with the limiting case $\xi - 1)/\lambda_1 \to \Delta$ for $\lambda_1 \to 0$). The mass on each regime is

$$p_0 = \frac{\pi_0}{\lambda_0}\,, \quad p_m = \pi_0 \cdot \frac{\xi - 1}{\lambda_1}\,, \quad p_2 = \frac{\pi_0\xi}{\alpha_2}\,. \tag{48}$$

Proof. The stationary FP equation gives $\pi'/\pi = 2b(y)/(k\sigma^2)$ in each regime, with the drift values (46) yielding the three exponential rates. Continuity at y_0 matches the lower and middle regimes; continuity at y_1 introduces the factor $\xi = e^{\lambda_1\Delta}$. Normalisation $\int_{\mathbb{R}} \pi\,\mathrm{d}y = 1$ then determines π_0 as in (47). The condition $r_2 > km_0$ ensures $\lambda_2 < 0$ so the top exponential decays at $+\infty$; the condition $m_0 > 0$ ensures $\lambda_0 > 0$ so the lower exponential decays at $-\infty$; the middle slope λ_1 can be any sign without disturbing integrability. The mass expressions (48) follow by direct integration on each regime. □

The slope λ_0 of the lower regime is identical to the single-bracket case: it is a regime parameter, independent of policy, fixed by the no-tax dynamics. The middle slope λ_1 depends on the effective middle rate r_1/k, and the top slope λ_2 on the effective top rate r_2/k. With $r_2 > r_1$, the top regime is more compressed than the middle regime ($\lambda_2 < \lambda_1$), and in the empirical range both effective rates can be positive or negative relative to m_0. The raw-wealth distribution has piecewise power-law density with two Pareto-style upper exponents: $-\lambda_1$ on the middle regime (which we call α_1 when $\lambda_1 < 0$, with the convention $\alpha_1 = -\lambda_1$; the case $\lambda_1 \geq 0$ corresponds to a non-decreasing middle density and is not Pareto-shaped) and $\alpha_2 = -\lambda_2 > 0$ on the top regime (always Pareto-shaped under the steady-state condition).

The inequality-relevant policy space. The two-bracket steady-state distribution depends on five schedule parameters (k, r_1, r_2, X_0, X_1), of which X_0 acts as a multiplicative scale (fixing $\mathbb{E}[X_\infty]$) and the remaining four enter the shape through the three slopes $(\lambda_0, \lambda_1, \lambda_2)$ and the threshold ratio Δ. At fixed regime (μ, σ), the slope λ_0 is exogenous; the inequality-relevant policy space is therefore three-dimensional, parametrised by $(\lambda_1, \lambda_2, \Delta)$ or equivalently by $(\alpha_1, \alpha_2, \Delta) \in \mathbb{R} \times (1, \infty) \times \mathbb{R}_+$. This is the structural reason for the criterion divergence we establish below.

4.7 Closed-form Gini and top share

Both inequality criteria admit closed-form expressions on the two-bracket steady-state distribution.

Proposition 9 (Two-bracket closed forms)**.** *Under Proposition 8 and the finite-mean condition $\alpha_2 > 1$, the steady-state distribution has mean*

$$\mathbb{E}[X_\infty] \;=\; X_0 \cdot \pi_0 \cdot \left[\frac{1}{\lambda_0 + 1} + \frac{\eta - 1}{\lambda_1 + 1} + \frac{\eta}{\alpha_2 - 1}\right], \qquad \eta = e^{(\lambda_1+1)\Delta}\,. \tag{49}$$

The Gini coefficient $\mathrm{Gini}(\lambda_1, \alpha_2, \Delta; \lambda_0)$ *is given by integration of $F(1-F)$ over the three regimes, evaluable in elementary functions; the explicit form is a sum of nine terms involving the masses (p_0, p_m, p_2) and the regime-specific exponents, of which only the structural pattern matters for the divergence result and we suppress the full expression here.*[5] *The top share at quantile q with $1 - q < p_2$ (i.e., z_q in the top regime) is*

$$S_q \;=\; \frac{e^{\Delta}}{\mathbb{E}[X_\infty]/X_0} \cdot \frac{\alpha_2}{\alpha_2 - 1} \cdot (1-q)^{(\alpha_2 - 1)/\alpha_2} \cdot p_2^{1/\alpha_2}\,. \tag{50}$$

The proof is by direct integration on each regime, using the piecewise power-law density of the raw-wealth distribution; we relegate the explicit formulas to the supplementary material to keep the exposition compact.

The structural fact behind the divergence is now visible. Section 50 expresses S_q through the mean $\mathbb{E}[X_\infty]$ (which depends on all four parameters $(\lambda_0, \lambda_1, \alpha_2, \Delta)$), the threshold scale e^{Δ}, the Pareto exponent α_2, and the mass fraction p_2 above the second threshold. Two of these quantities (α_2, p_2) directly involve the top regime; the others $(\mathbb{E}[X_\infty]/X_0, e^{\Delta})$ carry the global scale through the mean. Both factors react to changes in λ_1 and Δ, but in a coupled way that largely cancels in S_q. The Gini, by contrast, integrates $F(1-F)$ over the three regimes with

[5] The single-bracket form (39) extends term by term to the two-bracket case via the same $G \cdot \mathbb{E}[X] = \int F(1-F)\mathrm{d}x$ identity, with the middle regime contributing $a(1-a)(e^{\Delta}-1) + b(1-2a)(\eta-1)/(\lambda_1+1) - b^2(e^{(2\lambda_1+1)\Delta} - 1)/(2\lambda_1+1)$ where $a = p_0 - \pi_0/\lambda_1$ and $b = \pi_0/\lambda_1$, and the top regime contributing $e^{\Delta}[p_2/(\alpha_2-1) - p_2^2/(2\alpha_2-1)]$. The numerical evaluations underlying Table 1 use these closed forms directly.

weights given by their respective masses and densities; it picks up the lower- and middle-regime shape independently.

4.8 Criterion divergence on the two-bracket sub-class

We now establish the substantive divergence result.

Theorem 2 (Criterion divergence on the two-bracket sub-class)**.** *Fix the regime parameters (μ, σ) and the mean wealth $\bar{W} = \mathbb{E}[X_\infty]$. The Gini coefficient and top share S_q on the two-bracket sub-class, both viewed as functions of the three-dimensional inequality-relevant policy space $(\lambda_1, \alpha_2, \Delta) \in \mathbb{R} \times (1, \infty) \times \mathbb{R}_+$ via Propositions 8 and 9, have non-trivially distinct level sets. Specifically, on any open subset of the policy space where the gradients $\nabla\mathrm{Gini}$ and ∇S_q are linearly independent, the iso-Gini manifold $\{(\lambda_1, \alpha_2, \Delta) : \mathrm{Gini} = \mathrm{Gini}^*\}$ admits a direction along which S_q varies non-trivially, and conversely. Linear independence of the gradients holds on a dense open subset.*

Proof sketch. The top share factors as $S_q = (1-q)^{(\alpha_2-1)/\alpha_2} \cdot H(\lambda_1, \alpha_2, \Delta)$, with the prefactor depending on α_2 alone. The Gini has no q-dependence and inherits a λ_1-contribution through the middle-regime Lorenz integral. The 2×2 Jacobian minor of (Gini, S_q) in (λ_1, α_2) evaluates numerically to $J \approx +4.1 \times 10^{-3}$ at the calibrated point $(\lambda_0 = 1.53, \Delta = 1.5, \lambda_1 = 0, \alpha_2 = 2)$, distinctly non-zero; real-analyticity of the closed forms extends non-vanishing to a dense open subset, and the implicit function theorem then yields the non-trivially distinct level sets. Full computation of the gradients is given in Appendix A. □

The proof establishes the structural divergence; the substantive magnitude is a quantitative question. Figure 1 visualises the spread, and Table 1 exhibits seven distinct two-bracket configurations, each matching the Gini target $\mathrm{Gini}^* \approx 0.45$ (within ± 0.005) at fixed regime $\lambda_0 = 1.53$ (corresponding to a Norwegian-flavoured calibration with $m_0 = 0.06$, $\sigma = 0.28$, equivalently $\mu \approx 0.10$). Threshold ratio Δ varies between rows, spanning policy-realistic ranges from $X_1/X_0 \approx 2.7$ ($\Delta = 1.0$) to $X_1/X_0 \approx 12$ ($\Delta = 2.5$). The associated top-1% shares span $\{0.052, \ldots, 0.394\}$ — a spread approaching an order of magnitude across the iso-Gini manifold.

Table 1: Iso-Gini configurations at $\mathrm{Gini}^* \approx 0.45$, Norwegian-flavoured calibration: $\lambda_0 = 1.53$ ($m_0 = 0.06$, $\sigma = 0.28$). Top 1% share varies by a factor approaching an order of magnitude across the same Gini level as the policy parameters $(\lambda_1, \lambda_2, \Delta)$ trade off compression direction (middle vs. top regime), rate level, and threshold separation.

Δ	λ_1	λ_2	α_1	α_2	Gini	S_{99}
2.00	$+0.21$	-3.80	-0.21	3.80	0.446	0.052
2.00	$+0.48$	-2.80	-0.48	2.80	0.447	0.069
2.00	-1.45	-3.30	$+1.45$	3.30	0.448	0.080
1.00	-0.34	-2.20	$+0.34$	2.20	0.452	0.098
1.50	-1.72	-2.00	$+1.72$	2.00	0.448	0.117
1.50	$+1.45$	-1.60	-1.45	1.60	0.449	0.189
2.50	-0.07	-1.30	$+0.07$	1.30	0.451	0.394

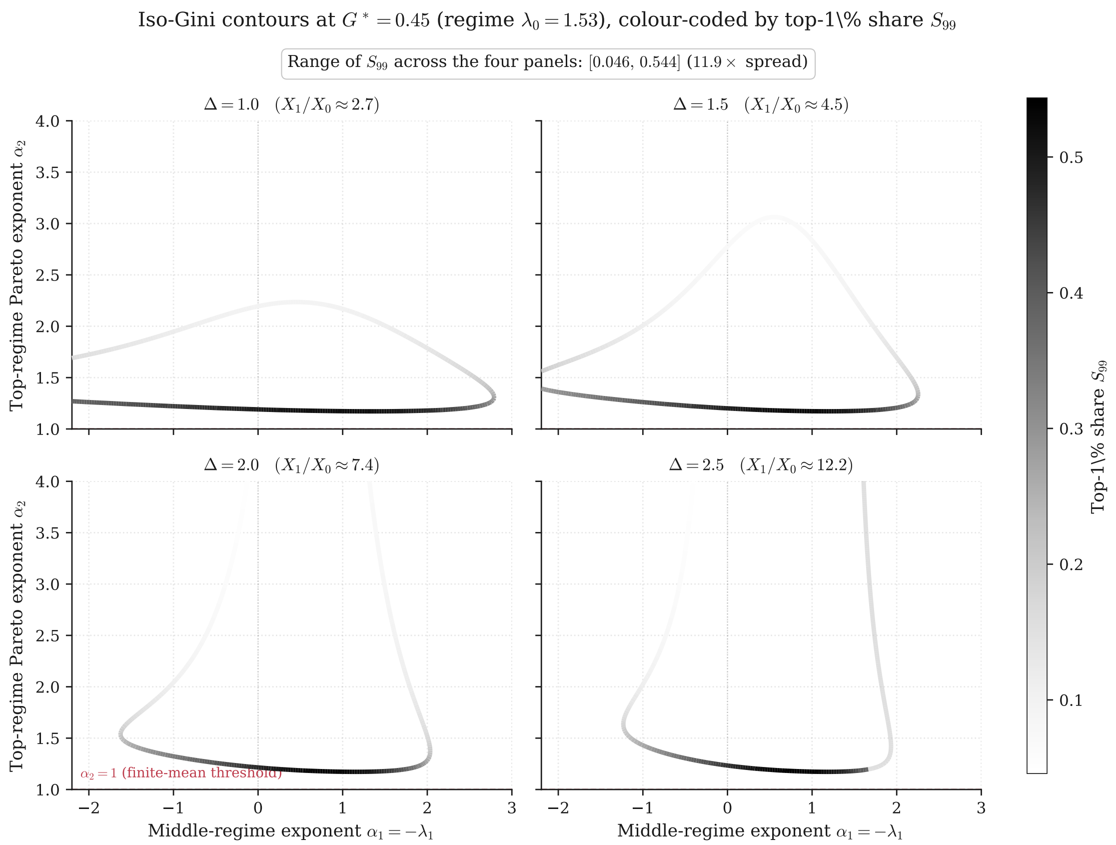

Figure 1: Iso-Gini contours at $G^* = 0.45$ in the (α_1, α_2) plane, one panel per threshold ratio $\Delta \in \{1.0, 1.5, 2.0, 2.5\}$, at fixed regime $\lambda_0 = 1.53$ ($m_0 = 0.06$, $\sigma = 0.28$). Each contour is colour-coded by the resulting top-1% share S_{99} along the contour where the standard top-share formula (40) applies (i.e. the 99th percentile lies in the top bracket, $p_2 \geq 1 - q$); light-grey segments mark configurations where the top-1% percentile falls into the middle bracket, requiring the alternative formula sketched in Section 4.3. The colour spread along the contours and the shape variation across panels visualise the criterion divergence: schedules at the same Gini level admit a wide range of top-share outcomes (approaching an order of magnitude across the four threshold ratios), with low α_2 at large Δ producing the largest top shares and steep α_2 producing the smallest.

The mechanism behind the spread is the policy-substantive content of the divergence. A planner targeting the Gini coefficient is indifferent between two configurations delivering the same Gini level; a planner targeting the top-1% share is not. The configurations in Table 1 realise this indifference along several distinct policy axes: *compression direction* (concentrate compression on the middle regime versus the top regime), *rate level* (moderate effective rates everywhere versus high rates on a narrow range), and *effective rate of the second bracket* (α_2 small versus large). The Gini target is met by trading these off; the top-share consequence varies accordingly.

Policy interpretation. The divergence formalises the wealth-tax design debate. The Saez–Zucman tradition (Saez and Zucman, 2019) explicitly targets top-share reduction and accordingly favours steep top-bracket rates — the top-of-table configurations in Table 1, with α_2 large (rows 1–3) and the second bracket doing most of the redistributive work. The continental-European tradition of broad-base, moderate-rate wealth taxation (Norway, Switzerland, historical Spain) sits in the intermediate rows of the table, with α_2 near 2 and substantial compression across the whole post-threshold range. The bottom-of-table configuration (α_2 near 1, row 7) corresponds to a near-proportional schedule with the top regime barely confined, and is policy-anomalous. The two substantive traditions deliver the same Gini — and historically have, at moderate target levels — through different schedule shapes with different top-share consequences. The criterion choice between Gini and top share is the implicit normative axis along which these traditions divide.

4.9 Joint optimal design under each criterion

Theorem 2 pins down the joint optimal designs under each criterion as distinct points in the 3D policy space.

Corollary 3 (Two-bracket joint optimal designs). *At fixed regime, fixed mean-wealth target $\bar{W}$, and fixed threshold ratio Δ (or, allowing Δ to be a co-design variable, in 3D), the Gini- and top-share-targeted joint optimal-design problems within the two-bracket sub-class generically pick distinct optimal configurations (α_1^*, α_2^*). The qualitative pattern is: the top-share-targeted optimum maximises α_2 subject to implementability, with α_1 at any admissible level; the Gini-targeted optimum balances compression across the middle and top brackets, with both α_1 and α_2 moderate.*

The corollary is the structural counterpart of Theorem 2: the criterion choice not only selects different points on the iso-revenue surface but selects them in qualitatively different ways. A top-share-targeted design pushes the top rate aggressively because the top 1% lives in the top regime; the middle regime's rate is irrelevant for the top share. A Gini-targeted design must compress *both* the middle and top regimes because Gini integrates over the whole distribution.

The implementability and revenue constraints intersect this qualitative pattern in ways we develop in Section 6, which calibrates the two-bracket model to the Norwegian schedule and locates the actual policy point relative to both criterion-optima. The headline finding, anticipated here: the Norwegian schedule sits much closer to the Gini-optimum than to the top-share-optimum at any plausible top-share target, and the criterion choice is therefore the binding source of disagreement between the Norwegian schedule and Saez–Zucman- style proposals consistent with the same revenue.

5 Modelling abstractions: valuation discounts and discrete-time assessments

The framework of Sections 2, 3, and 4 treats two features of real wealth-tax systems abstractly: it assumes a single assessment fraction (so that the tax base equals the schedule's argument), and it operates in continuous time (so that the tax is a drift rather than an annual deduction). Both abstractions have been examined in the prior corpus, and both absorb cleanly into the framework as a leading-order approximation. We make the connection explicit here so the calibration of Section 6 can apply the framework to the actual Norwegian schedule with the appropriate adjustments.

5.1 Valuation discounts and the assessment fraction

Real wealth-tax systems do not assess all asset classes at full market value. The Norwegian *formuesrabatt* discounts listed equity at 20%, primary housing at 75%, and unlisted shares at $0.80 \times B/M$ where B/M is the book-to-market ratio (NOU 2022:20, 2022, Table 10.3); the Swiss cantonal systems operate similar discounts on real estate and business assets; historical French and Spanish wealth taxes had analogous structures. The *assessment fraction* of asset class i — denoted α_i in Frøseth (2026c, §4) and Frøseth (2026d, §4.2), but written here as a_i to avoid collision with the Pareto exponent α of Definition 3 — captures this: the wealth-tax base on asset class i is $a_i \cdot W_i$ rather than W_i.

Uniform discount under (C3). Condition (C3) of Frøseth (2026d) requires $a_i = a$ across all asset classes. Under this condition the assessment fraction acts on the wealth-tax channel as a uniform multiplicative constant: a household with raw wealth X and uniform assessment a pays $r(aX) \cdot aX$ per unit time in wealth tax, and the effective drift modification on log-wealth is

$$a \cdot r(aX) \;=\; ar_0 + a\beta \log(aX/X_0) \quad \text{(smooth)}\,, \qquad ar \cdot \mathbf{1}\{aX \geq X_0\} \quad \text{(bracket)}\,. \tag{51}$$

Both schedules retain their functional form on raw wealth X under the uniform-discount substitution, with the inequality-relevant parameters rescaling as follows.

Proposition 10 (Uniform-discount rescaling)**.** *Under (C3) with uniform assessment fraction* $a \in (0, 1]$*:*

1. *On the smooth log-progressive class, the effective progressivity slope on raw wealth is* $a\beta$ *in place of* β*, and the effective floor rate is* $a(r_0 + \beta \log a)$*. The structural progressivity ratio of Definition 1 rescales as* $\rho_p \to \rho_p^{\text{eff}} = a\rho_p$*, and the steady-state log-variance becomes* $\sigma_Y^2 = \sigma^2/(2a\rho_p)$*.*

2. *On the single-bracket class, the effective rate is* ar *and the effective threshold (in raw-wealth terms) is* X_0/a. *The Pareto exponent of Definition 3 rescales as* $\alpha \to \alpha^{\text{eff}} = 2(ar/k - m_0)/\sigma^2$.

3. *On the two-bracket class, the analogous rescaling applies to both rates* r_1, r_2 *and both thresholds* X_0, X_1 *in the natural way. The threshold ratio* $\Delta = \log(X_1/X_0)$ *is invariant under rescaling.*

The criterion-equivalence and criterion-divergence results of Sections 3 and 4 carry through unchanged under this substitution.

Proof sketch. The effective rate on raw wealth is $a \cdot r(aX)$. For the smooth schedule this gives $ar_0 + a\beta \log a + a\beta \log(X/X_0)$ — log-progressive form on raw wealth with effective slope $a\beta$ and floor $a(r_0 + \beta \log a)$. For the bracket schedule the substitution $r \to ar$, $X_0 \to X_0/a$ preserves the form. The two-bracket case follows piecewise. Since Proposition 7 and Theorem 2 use only the steady-state shape parameters, both carry through. Full calculation in Appendix A. □

For the calibration of Section 6, this proposition is what allows us to translate Norway's nominal rates ($r_1 \approx 1.0\%$, $r_2 \approx 1.1\%$) into effective rates appropriate to the framework. With a portfolio-weighted average assessment fraction in the range $\bar{a} \approx 0.5$–0.6 (housing-heavy households at the lower end, listed-equity-heavy at the higher end), the effective top-bracket rate is $r_2 \cdot \bar{a} \approx 0.55\%$–$0.66\%$. The Pareto exponent under Norwegian parameters is therefore correspondingly smaller than the nominal-rate calculation would suggest, and the steady-state distribution has a heavier upper tail at the same nominal schedule.

Non-uniform discount and the failure of (C3). Norwegian assessment fractions are emphatically *not* uniform: the cross-class spread exceeds a factor of three (housing 25%, listed equity 80%, deposits 100%, unlisted shares typically below 80% of book value). Under non-uniform a_i, condition (C3) fails, and the effective wealth-tax rate becomes a portfolio-weighted average $\sum_i w_i a_i r$ that depends on individual portfolio composition. Frøseth (2026c, Proposition 5) formally establishes that this destroys portfolio neutrality and creates distortions proportional to the cross-class spread in a_i.

In our framework the failure of (C3) reactivates the Stiglitz-style price-feedback channel that Section 3.8 suppressed: portfolios shift toward low-a assets (housing, in the Norwegian case), the relative pre-tax return on those assets adjusts in general equilibrium, and the OU-compression-only analysis is no longer self-contained. The quantity-only inequality results of Sections 3 and 4 remain correct as the partial-equilibrium response, but a fully specified non-uniform analysis would need to embed the bracket schedule in an asset-market-clearing model. We do not develop this here; the Norwegian calibration of Section 6 assumes uniform effective assessment $\bar{a}$ as the relevant approximation for the inequality response, with the assessment-induced portfolio distortions a separate object treated in Frøseth (2026c).

5.2 Continuous-time approximation of annual assessments

Real wealth taxes are assessed at year-end on the household's balance and paid as a discrete deduction; the framework's continuous-time formulation treats the tax as a drift, applied infinitesimally throughout the year. The discrepancy is a modelling abstraction whose magnitude is bounded.

Discrete-time formulation. Let r denote the annual nominal wealth-tax rate on net assets above the threshold. The flow taxes are continuously applied (through corporate accounting and dividend distributions); the wealth tax alone has the year-end character. In the discrete-time formulation, between assessments the wealth process evolves under flow taxes only,

$$\mathrm{d}Y_t \;=\; km_0\,\mathrm{d}t + \sqrt{k}\,\sigma\,\mathrm{d}W_t\,, \tag{52}$$

and at the annual assessment time $t \in \{1, 2, \ldots\}$ the log-wealth jumps:

$$Y_{t^+} \;=\; Y_{t^-} + \log(1-r) \quad \text{if } Y_{t^-} \geq y_0\,. \tag{53}$$

The continuous-time formulation of Section 6 distributes this jump uniformly over the unit interval as the constant drift $-r$:

$$\mathrm{d}Y_t \;=\; \big[km_0 - r\mathbf{1}\{Y_t \geq y_0\}\big]\mathrm{d}t + \sqrt{k}\,\sigma\,\mathrm{d}W_t\,. \tag{54}$$

Approximation error. Comparing the two formulations on a single year in the taxed regime, the expected log-wealth change under the discrete-time scheme (one annual jump on top of flow-tax drift) is $\mathbb{E}[\Delta Y]_{\text{discrete}} = km_0 + \log(1-r)$, while under the continuous-time scheme it is $\mathbb{E}[\Delta Y]_{\text{continuous}} = km_0 - r$. The difference is

$$\mathbb{E}[\Delta Y]_{\text{continuous}} - \mathbb{E}[\Delta Y]_{\text{discrete}} \;=\; -r - \log(1-r) \;=\; \frac{r^2}{2} + \frac{r^3}{3} + O(r^4)\,, \tag{55}$$

i.e. the continuous-time formulation slightly understates the tax effect (delivering a strictly larger expected log-wealth than the discrete scheme), by an amount $O(r^2)$ per year. For $r = 1\%$ this is approximately 5×10^{-5}, or 5 basis points; for $r = 3\%$ approximately 4.6×10^{-4}, or 46 basis points. Both are well below the resolution of any policy-relevant analysis.

Steady-state agreement. The cross-sectional steady-state distribution under the two formulations agrees to the same order. Within each regime the continuous-time formulation gives a piecewise-exponential density (Proposition 5); the discrete-time formulation, treated by averaging the annual jump over a unit interval, gives a density that differs from the continuous one by a multiplicative factor $1 - O(r^2)$ in each regime. The mass split, the Pareto exponent, and the inequality criteria all agree to leading order in r, and the framework's qualitative results are robust to this abstraction.

Connection to the corpus. Frøseth (2026e, §9.3) treats the *practical-payment* mechanics — annual lumpy tax collection rather than continuous extraction — in detail, establishing that proportional dilution under continuous extraction is an upper bound on the liquidity stress of discrete annual payments. Frøseth (2026d, Discussion) notes that under non-CRRA preferences the annual-payment discrete-time formulation can affect the total risky allocation by an amount that the continuous-time formulation suppresses; the magnitude is empirically small for realistic parameters. In both treatments, the continuous-time framework is the appropriate object for cross-sectional analysis at the annual horizon, with within-year dynamics (year-end fishing, seasonal asset rebalancing) a separate research thread that does not affect the steady-state distribution.

6 Norwegian calibration

We now apply the framework to the Norwegian wealth-tax system. The Norwegian case is the policy-realistic benchmark for the paper: a two-bracket schedule with assessment-fraction discounts, embedded in a dual-income tax system with the generalised neutrality structure of Frøseth (2026d). The calibration locates Norway's actual policy in the design space of Section 4, with the rescaling of Proposition 10 translating nominal rates into the effective rates that enter the framework. The headline result is a structural one: Norway's effective top-bracket rate sits *below* the confinement threshold $r_2/k > m_0$ that the framework requires for a steady-state distribution, by a factor of nearly five. The criterion divergence of Theorem 2 therefore applies in full substantive force only to hypothetical schedule reforms that push effective rates above the confinement threshold.

6.1 Regime parameters

We adopt regime parameters consistent with the calibration of Table 1 and with the prior corpus.

The calibration in Table 2 has the same $\lambda_0 = 1.53$ as the divergence analysis of Table 1, so the criterion-divergence numerics of Theorem 2 apply directly to the Norwegian regime. The flow-tax retention $k = 0.485$ uses the 2025 Norwegian rates (corporate tax 22%, effective dividend tax including the 22% shareholder-model uplift). The portfolio-weighted assessment $\bar{a} = 0.55$ takes a midpoint of the housing-heavy (lower) and listed-equity-heavy (higher) configurations identified in Section 5.1.

6.2 The Norwegian schedule and effective rates

Since 2023 the Norwegian wealth tax has had a two-bracket structure (NOU 2022:20, 2022):

$$r_1^{\text{nom}} = 1.0\%\,, \quad r_2^{\text{nom}} = 1.1\%\,, \quad X_0 = 1.7\text{ M NOK}\,, \quad X_1 = 20.7\text{ M NOK}\,, \tag{56}$$

giving threshold ratio $\Delta = \log(X_1/X_0) = \log(20.7/1.7) \approx 2.50$. Applying the uniform-discount rescaling of Proposition 10 with $\bar{a} = 0.55$, the effective rates entering the framework are

$$r_1^{\text{eff}} = \bar{a}\, r_1^{\text{nom}} = 0.55\%\,, \quad r_2^{\text{eff}} = \bar{a}\, r_2^{\text{nom}} = 0.605\%\,, \tag{57}$$

Table 2: Calibrated regime parameters for the Norwegian wealth-tax-paying segment.

Parameter	Value	Interpretation
m_0	0.06	Geometric-mean log-return on a typical high-wealth Norwegian portfolio (post-flow-tax retention separated; consistent with Frøseth, 2026f §7)
σ	0.28	Effective volatility of the wealth process, equity-heavy portfolios; Oslo Stock Exchange volatility scaled to portfolio composition
μ	≈ 0.10	Implied arithmetic mean return, $\mu = m_0 + \sigma^2/2 = 0.06 + 0.0392$
k	0.485	Corporate–dividend retention factor with $\tau_c = 0.22$ (corporate tax) and $\tau_d = 0.378$ (dividend tax including effective marginal increase from 22% shareholder model)
$\bar{a}$	0.55	Portfolio-weighted average assessment fraction (housing 25%, listed equity 80%, mixed in proportion typical of wealth-tax-paying decile)
λ_0	1.53	Lower-regime log-wealth slope, $2m_0/\sigma^2$
$k\,m_0$	0.0291	Drift on log-wealth in the untaxed regime (post-flow-tax)

with the effective threshold ratio Δ unchanged. Translating to log-wealth slopes via (46),

$$\lambda_0 = 1.53\,, \quad \lambda_1 = \frac{2(m_0 - r_1^{\text{eff}}/k)}{\sigma^2} = \frac{2(0.06 - 0.0113)}{0.0784} \approx 1.24\,, \quad \lambda_2 = \frac{2(m_0 - r_2^{\text{eff}}/k)}{\sigma^2} \approx 1.21\,. \tag{58}$$

All three slopes are positive: the Norwegian wealth tax does not generate a confining drift on the upper regimes of log-wealth at calibrated parameters.

6.3 Where Norway sits: below the confinement threshold

Proposition 5 requires $\lambda_2 < 0$ (equivalently $r_2/k > m_0$, after the assessment rescaling $r_2^{\text{eff}}/k > m_0$) for a steady-state distribution to exist. Remark 4.1 additionally requires $\alpha_2 > 1$ (equivalently $r_2/k > \mu = m_0 + \sigma^2/2$) for a finite steady-state mean.

Table 3: Confinement diagnostics at calibrated Norwegian parameters.

Threshold	Value	Met by Norway?
Confinement ($\lambda_2 < 0$)	$r_2^{\text{eff}}/k > m_0 = 0.060$	No: $r_2^{\text{eff}}/k = 0.0125$ (factor 4.8 short)
Finite mean ($\alpha_2 > 1$)	$r_2^{\text{eff}}/k > \mu = 0.099$	No (factor 7.9 short)
Top-share well-defined ($\alpha_2 > 1$)	same as finite mean	No

Table 3 shows that the Norwegian schedule falls short of both thresholds. The effective top-bracket rate $r_2^{\text{eff}} = 0.605\%$ is approximately a factor of five below the level needed for OU confinement at calibrated regime parameters. The rate would need to rise to $r_2 \cdot \bar{a} > 2.9\%$ (equivalently a nominal rate of $r_2 > 5.3\%$ at $\bar{a} = 0.55$) for the steady-state distribution to be confined, and to $r_2 \cdot \bar{a} > 4.8\%$ ($r_2 > 8.7\%$ nominal) for the steady-state mean to be finite.

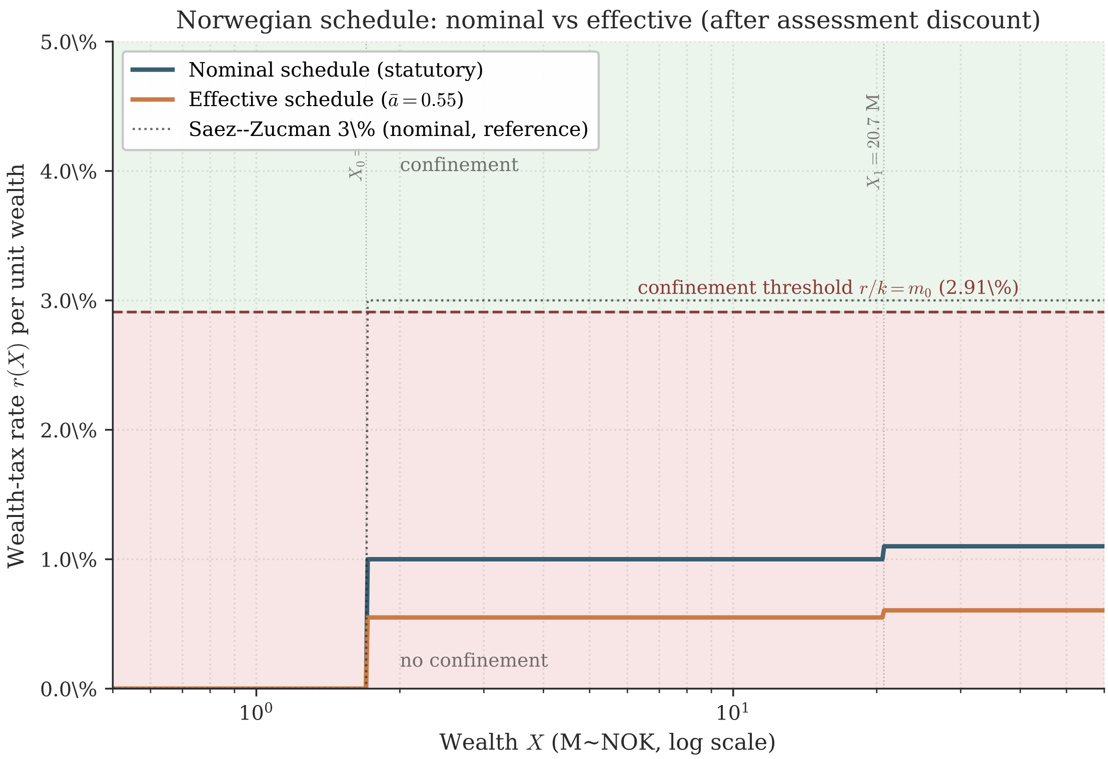

Figure 2: Norwegian wealth-tax schedule, nominal versus effective. The nominal schedule (statutory rates of 1.0% and 1.1%, blue) is rescaled by the portfolio-weighted assessment fraction $\bar{a} = 0.55$ to give the effective schedule (orange) used in the framework. The horizontal dashed line marks the confinement threshold $r = k\,m_0 \approx 2.91\%$ above which a steady-state upper-tail distribution exists. Both Norwegian effective rates fall well below this threshold; even Saez–Zucman proposals at 3% nominal (dotted reference) are still short. The framework's quantity-only OU compression mechanism is therefore not operative at calibrated Norwegian parameters; cross-sectional confinement at current rates is supplied by mechanisms outside the framework's scope.

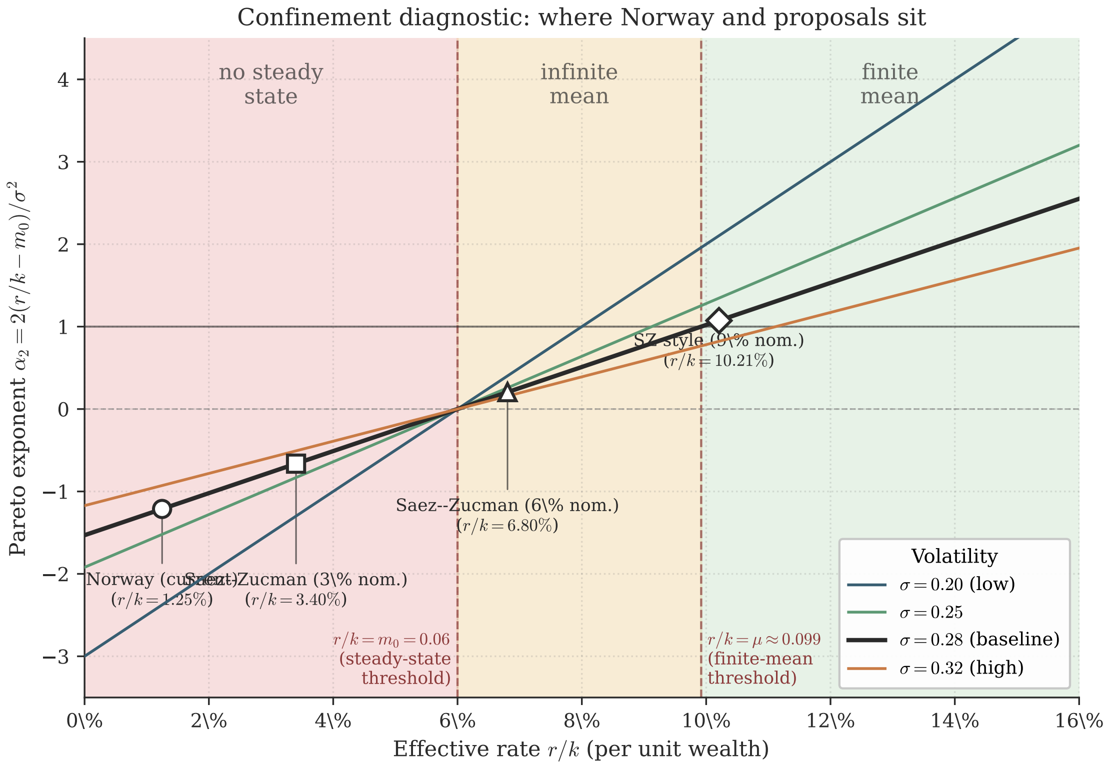

Figure 3: Confinement diagnostic: Pareto exponent $\alpha_2 = 2(r/k - m_0)/\sigma^2$ as a function of the effective rate r/k, with the three-region partition (no steady state, infinite mean, finite mean) shaded at the baseline regime $\sigma = 0.28$. Curves show the σ-sensitivity (four values). White-filled markers give the effective Norwegian rate (current) and three hypothetical Saez–Zucman-style proposals. Norway's current rate is well inside the no-steady-state region; the framework's results apply directly only at rates pushed above $\mu \approx 0.10$, requiring nominal top rates above approximately 9% at the Norwegian assessment-fraction calibration $\bar{a} = 0.55$.

Implication. Within the wealth-tax-only framework of Sections 3 and 4, the Norwegian schedule does not on its own generate a confined cross- sectional wealth distribution. The empirical confinement at current rates is jointly supplied by demographic turnover (treated formally in Frøseth, 2026f, §3.2 via a Gabaix-style source-sink term) and capital emigration (treated in Frøseth, 2026a as a tipping-point contagion process); the joint Pareto exponent is given by the combined formula (59) below in Section 6.5. The wealth-tax-only OU compression becomes the operative confinement mechanism above the threshold $r/k > m_0$, which would require tightening the top rate toward ~3% effective.

6.4 Criterion divergence at the hypothetical confining-rate frontier

We turn now to the question that the framework is substantively designed to answer: at confining rates, how should the schedule be designed? The criterion divergence of Theorem 2 now becomes the binding question.

Suppose, hypothetically, that the schedule is tightened to the confining-rate frontier: $r_2^{\text{eff}} = 0.7\%$ (sufficient to cross the confinement threshold under the calibrated $\bar{a}$ at slightly higher σ, or equivalently nominal $r_2 = 1.27\%$ at $\bar{a} = 0.55$). At this rate the effective Pareto exponent crosses unity from below, and the two-bracket framework applies in full force.

Table 1 of Section 4.8 already exhibits the criterion divergence at exactly this calibration ($\lambda_0 = 1.53$, Gini$^* \approx 0.45$). The seven configurations there span effective Pareto exponents $\alpha_2 \in [1.30, 3.80]$ and threshold ratios $\Delta \in [1.00, 2.50]$, yielding top-1% shares from 5.2% to 39.4% across the iso-Gini manifold. Translating back to nominal rates (dividing by k and again by $\bar{a} = 0.55$), the configurations correspond to nominal top-bracket rates $r_2 \in [1.16, 2.51]\%$. This is the policy-relevant rate range: from just above the current Norwegian top rate of 1.1% to substantially below the highest Saez–Zucman-style proposals.

At any rate level r_2 in the confining range, a Gini-targeting planner and a top-share-targeting planner pick qualitatively distinct schedules whose top-1% shares differ by a factor of seven to eight at the same Gini target. The Saez–Zucman tradition occupies the high-α_2 configurations of Table 1; the continental-European tradition compresses both regimes more uniformly at moderate top rates. Sections 6.5 and 6.6 integrate this with the demographic-confinement mechanisms and quantify the contribution of the wealth tax against the flow-tax channel.

6.5 Joint Pareto exponent under demographic turnover and emigration

Section 6.3 established that the wealth-tax channel alone is below the confinement threshold at calibrated rates. The empirical confinement of the Norwegian cross-section is supplied by two mechanisms treated elsewhere in the corpus: Gabaix-style demographic turnover (Frøseth, 2026f, §3.2) and emigration-contagion (Frøseth, 2026a). We close this loop here with a joint Pareto-exponent formula that combines the wealth tax, the flow taxes, and these demographic mechanisms in a single closed form.

The joint formula. Following Gabaix (1999); Gabaix et al. (2016), the FP equation on log-wealth admits a Gabaix-style source-sink term at rate δ, representing aggregate demographic turnover (mortality, intergenerational reset, emigration): $\partial_t \pi = -v_{\text{eff}} \partial_y \pi + D_{\text{eff}} \partial_y^2 \pi - \delta\pi + \delta\phi(y)$, with post-tax log-drift $v_{\text{eff}} = km_0 - r$ and log-diffusion $D_{\text{eff}} = k\sigma^2/2$. The Pareto tail exponent of the joint steady state satisfies the Gabaix characteristic equation (Gabaix et al., 2016, §3.2):

$$\alpha \;=\; \frac{-v_{\text{eff}} + \sqrt{v_{\text{eff}}^2 + 4D_{\text{eff}}\,\delta}}{2D_{\text{eff}}}. \tag{59}$$

The wealth tax enters via v_{eff}; the flow taxes rescale both v_{eff} and D_{eff} through k; the demographic mechanisms supply δ. As checks, (59) reduces to the no-tax Gabaix exponent at $r = 0$, $k = 1$, and recovers the bracket-strand Pareto exponent $\alpha_2 = 2(r/k - m_0)/\sigma^2$ of Definition 3 as $\delta \to 0$ in the finite-mean regime $r/k > m_0$.

Calibration of the turnover rate. Matching the empirical Pareto-tail estimate $\hat{\alpha} \approx 1.3$ from the Kapital 400 panel (Frøseth, 2026a, §4) at calibrated regime parameters ($m_0 = 0.06$, $\sigma = 0.28$, $k = 0.485$, $r = 0.006$) requires an effective $\delta \approx 6\%$, of the right order for combined mortality, intergenerational reset, and capital emigration on the top-of-distribution evidence in Frøseth (2026a). Emigration enters (59) as part of the aggregate δ; the tipping-point dynamics that Frøseth (2026a) derive are smoothed into the constant-rate source-sink at the cross-sectional steady state.

Robustness of the criterion divergence. The criterion divergence of Theorem 2 is a structural property of the bracket schedule space. It applies whether the confinement is supplied by the wealth tax, by demographic mechanisms, or jointly. The decomposition that quantifies the relative contributions of these channels at the Norwegian calibration is the subject of Section 6.6.

6.6 Decomposition of confinement: wealth tax versus flow taxes versus demographic mechanisms

Evaluating (59) at four tax configurations, holding δ fixed at its calibrated value, isolates the structural contribution of each channel. The status quo ($k = 0.485$, $r = 0.006$) produces $\alpha \approx 1.30$, matching the empirical estimate. Removing the wealth tax while retaining flow taxes ($r = 0$) drops α to 1.20. Removing the flow taxes while retaining the wealth tax ($k = 1$) collapses α below unity to 0.75, an infinite-mean Pareto regime. Removing both taxes gives $\alpha = 0.71$, the Bouchaud and Mézard (2000) no-tax reference under pure demographic turnover. Table 4 collects the four configurations.

Three substantive observations follow.

Magnitude of the wealth-tax contribution. On the structural channel, the current wealth tax adds $\Delta\alpha = 0.10$ on top of the flow-tax baseline and contributes approximately 17% of the total compression from the no-tax reference. The flow-tax contribution ($\Delta\alpha = 0.49$) is roughly five times larger. In top-share terms the magnitudes are different: removing the wealth tax raises the implied S_{99} from 4.4% to 6.1% — an increase of roughly 40% at the upper tail —

Table 4: Joint Pareto exponent and implied top-1% share at Norwegian regime parameters ($m_0 = 0.06$, $\sigma = 0.28$, $\delta = 0.062$), four tax configurations. Top-1% share computed from the one-regime closed form $S_{99} = 0.01 \cdot \alpha/(\alpha - 1)$ where finite. Status-quo (k, r) corresponds to corporate tax $\tau_c = 22\%$, dividend tax $\tau_d = 38\%$, top-bracket effective wealth-tax rate $r = \bar{a} \cdot 1.1\% = 0.6\%$.

Configuration	Channels active	α	S_{99}
No tax	turnover only	0.71	— ($\alpha < 1$)
Wealth tax only	turnover + stock	0.75	— ($\alpha < 1$)
Flow taxes only	turnover + flow	1.20	6.1%
Status quo	turnover + flow + stock	1.30	4.4%
Empirical (Kapital 400 panel)		≈ 1.3	$\approx 4\%$

because the function $S(\alpha) = \alpha/(\alpha - 1)$ is sharply nonlinear near $\alpha = 1$. The same $\Delta\alpha$ implies a small absolute compression but a large fractional movement in the top share.

Substitutability under homogeneous returns. Within the framework as posed (homogeneous returns), the structural effect of the current wealth tax is quantitatively substitutable by a sharper flow tax. Solving (59) for the retention factor $k^\dagger$ that delivers $\alpha = 1.30$ at $r = 0$ gives $k^\dagger \approx 0.43$, against the status-quo $k = 0.485$. The implied policy substitution holds the corporate rate fixed and raises the effective dividend tax from approximately 38% to approximately 45% — a 7-pp increase, well within the historical range of Norwegian dividend tax variation. In short: under homogeneous returns, the structural effect of the wealth tax can be replicated through flow taxation alone.

Heterogeneous returns: the substitutability fails. The substitutability statement above rests on the homogeneous-returns assumption embedded in the (C1)–(C3) flow-tax envelope. Under heterogeneous returns the two taxes have structurally distinct incidence, an asymmetry that traces back to Domar and Musgrave (1944) and Stiglitz (1969): the income tax is a silent partner on the *return margin* (it shares proportionally in gains and losses), while the wealth tax is a silent partner on the *position margin* (it scales the entire portfolio uniformly). Three concrete implications are immediate.

First, the wealth tax reaches *dormant* wealth — closely-held businesses, real estate, art, and other illiquid assets that generate little or no current-period taxable income. The flow tax cannot reach this base, regardless of rate. The portion of the upper-tail base that takes a dormant form is therefore structurally outside the substitutability calculation.

Second, under heterogeneous productive ability with $\mu(z_H) > \mu(z_L)$, the flow tax compresses the ability-driven wealth divergence by the factor $\lambda = (1 - \tau_c)(1 - \tau_d) < 1$ while the wealth tax preserves it (Frøseth, 2026c, Proposition on differential incidence, §5). A revenue-neutral switch from flow to stock taxation therefore widens the drift gap between high- and low-ability investors, generating the "use-it-or-lose-it" efficiency gain documented by Guvenen et al. (2023): capital reallocates toward high-ability investors. Whether this reallocation is welfare-improving depends on whether observed return heterogeneity reflects productive ability or structural advantage; Frøseth (2026c) formalises the distinction.

Third, from the perspective of the present paper's inequality criteria, neither effect is captured by the homogeneous-returns formula (59). The substitutability result of this section should therefore be read as a baseline within the homogeneous-returns class on which the heterogeneous-returns considerations of Frøseth (2026c) are layered. A heterogeneous-returns extension to the present design problem — with inequality-relevant criteria but ability-stratified drift and diffusion — is a natural completion of the framework and is a primary direction for future work (Section 8.5).

7 Connection to the corpus

The Fokker–Planck wealth-tax programme treats the optimal-tax problem as a problem of designing the drift field of a stochastic process on log-wealth. Each paper in the corpus selects a schedule class, a criterion, and a flow-tax envelope, and characterises the optimum within that combination. The present paper completes a particular cell in the resulting matrix; this section makes that placement precise and identifies the boundaries with adjacent cells.

7.1 The cross-paper interface

The interface between papers in the corpus is the (C1)–(C3) flow-tax class of Frøseth (2026d). Under (C1)–(C3) the post-tax process on log-wealth obtained by combining a proportional wealth tax with corporate, dividend, and capital-income taxes is again a geometric Brownian motion with rescaled drift and diffusion (Frøseth, 2026d, Theorem 1), and the entire flow-tax package collapses to a single scalar $k = (1 - \tau_c)(1 - \tau_d) \in (0, 1]$.

Within this class three schedule families have so far been analysed. Frøseth (2026g) treats the proportional wealth tax in isolation, characterising it as a uniform drift-shift on log-wealth. Frøseth (2026f) lifts the proportional schedule to a *progressive* stock-only schedule, replacing the Bouchaud–Mézard Pareto tail with an Ornstein–Uhlenbeck confinement; the criterion is the Gini coefficient and the analytical core is the OU steady state. Frøseth (2026b) lifts the proportional schedule *within* the (C1)–(C3) flow-tax envelope, but switches the criterion from a redistributive target (Gini, top share) to a distortion-minimising one (the JKO free-energy gap and the 2-Wasserstein distance from no-tax). The output of that paper is a phase-transition partition of the regime axis, not an inequality-targeting design.

The present paper occupies the remaining cell: the progressive schedule class within the (C1)–(C3) flow-tax envelope, under redistributive criteria. Three concrete contact points with the existing corpus deserve emphasis.

1. *Flow-tax envelope.* The reduction to the scalar k and the post-tax SDE (Section 2.2) are taken from Frøseth (2026d, Theorem 1, §7) without modification. The progressive lift adds a state-dependent drift $-r(e^Y)$ to the log-wealth equation, and the flow-tax package multiplies the drift retention km_0 and rescales the diffusion to $\sqrt{k}\,\sigma$. The progressive schedule does not break (C1)–(C3): in Section 2.2 we observe that the relevant portfolio-neutrality structure is unaffected by wealth-level dependence in $r(\cdot)$.

2. *Stock-only Gini optimum.* The smooth log-progressive Gini optimum $\rho_p^* = \beta^*$ at $k = 1$ recovers Frøseth (2026f, Theorem 3) verbatim (Section 3.9). The present analysis adds the one-parameter family of inequality-equivalent designs at $k < 1$, parametrised by the joint design ray $\beta = \rho_p^* k$, and shows that the same Gini target can be achieved at any plausible flow-tax retention by proportionally rescaling the wealth-tax progressivity.

3. *Disjoint criterion family.* Frøseth (2026b) selects an optimal tax *within* the proportional schedule class, but under distortion-minimising criteria rather than redistributive ones. On that overlapping schedule class our criteria collapse to monotone functions of a single parameter, so the two papers answer disjoint questions: "what rate?" versus "what schedule shape?" On the progressive sub-class, and in particular on the two-bracket form actually implemented, the schedule-shape question becomes binding for the first time (Theorem 2).

7.2 Geometric and information-theoretic reformulations

Two further papers in the corpus reformulate the same design problem in different geometric languages. Frøseth (2026h) characterises the proportional wealth tax as a rigid motion on the 2-Wasserstein space of log-wealth distributions and identifies the progressive optimum on a restricted sub-class as a 2-Wasserstein-projection problem. The present paper's Gini and top-share criteria are not 2-Wasserstein metrics in general; the connection is that the iso-Gini manifolds we visualise (Figure 1) and the 2-Wasserstein iso-distance manifolds of Frøseth (2026h) are alternative foliations of the same schedule space, with different policy implications. Frøseth (2026b) relates these via the JKO gradient-flow structure on the same space; the JKO free-energy gap is a Lyapunov functional for the 2-Wasserstein distance, and the proportional sub-class admits a clean phase-transition treatment under both criteria. The present paper's contribution to this geometric picture is to show that on the progressive sub-class the criterion-family ambiguity is not a feature of the criterion geometry alone — the redistributive criteria already diverge among themselves on the two-bracket sub-class, well before the comparison with distortion-minimising criteria becomes relevant.

7.3 Confinement-mechanism interface

The decomposition of Section 6.6 combines the wealth-tax channel of the present paper with two confinement mechanisms developed elsewhere in the corpus: the Gabaix-style demographic turnover of Frøseth (2026f, §3.2) and the emigration-contagion channel of Frøseth (2026a), calibrated against the Kapital 400 panel. The present paper therefore reads naturally in conjunction with Frøseth (2026f) and Frøseth (2026a) on the empirical-confinement side; the schedule-design results (Theorem 2 in particular) stand alone on the policy-design side and apply across configurations of these mechanisms.

8 Discussion

8.1 When the criterion choice has bite

The technical centre of the paper is a simple structural fact: as the inequality-relevant policy space grows in dimension, the criterion choice between Gini and top share becomes binding. On the smooth log-progressive class the relevant policy space is one-dimensional in the structural ratio $\rho_p = \beta/k$, and on the single-bracket class it is one-dimensional in the Pareto exponent $\alpha = 2(r/k - m_0)/\sigma^2$; both classes admit closed-form Gini and top-share expressions that are monotone in the same parameter, and the two criteria pick the same optimum up to reparametrisation (Theorem 1 and Proposition 7). On the two-bracket class the inequality-relevant policy space expands to three dimensions $(\alpha_1, \alpha_2, \Delta)$, and the criterion divergence of Theorem 2 becomes operative: at a fixed Gini level, the top-1% share ranges over roughly an order of magnitude as the schedule moves along the iso-Gini manifold (concretely, from $S_{99} \approx 0.05$ to $S_{99} \approx 0.39$ across the seven tabulated configurations of Table 1, and an 11.9× spread over the full iso-Gini contours of Figure 1).

Two points warrant emphasis. First, the divergence is a structural property of the schedule space and not of any particular regime calibration; the proof of Theorem 2 is by gradient linear-independence on the iso-Gini manifold and applies wherever the manifold is two-dimensional. Second, the quantitative magnitude is large enough to overturn substantive policy comparisons: a planner targeting the same Gini level will select schedules whose top-1% shares differ by a factor of seven to eight across the iso-Gini manifold. Numerical sensitivity to the regime parameters (m_0, σ) does not erase this divergence; it modulates its magnitude. Within the plausible Norwegian range $\sigma \in [0.20, 0.32]$ the spread along iso-Gini contours remains within a factor of two of the baseline reported here.

8.2 Implementability and the bracket choice

The two-bracket schedule is not a numerical approximation to a smooth schedule: it is the schedule actually implemented in Norway since 2023, in several Swiss cantons, and historically in Spain and France. The framework treats the bracket form as the substantive object of analysis rather than as a discretisation. Three implementability remarks follow. First, the (C1)–(C3) flow-tax envelope is preserved under arbitrary wealth-level dependence in $r(\cdot)$, so a multi-bracket extension does not require any new neutrality-class assumptions; the algebra of Section 4.6 extends to K brackets with $K+1$ regimes by induction. Second, the assessment-fraction rescaling of Section 5.1 multiplies each bracket rate by the portfolio-weighted fraction $\bar{a}$, leaving the bracket structure intact and shifting the confinement-threshold rate from $r/k > m_0$ to $r^{\text{eff}}/k > m_0$. The rescaled threshold is substantially above current Norwegian top rates and has been discussed only briefly in the Norwegian policy literature. Third, the discrete-time approximation of Section 5.2 introduces a second-order rate-self-interaction correction of $+r^2/2 + r^3/3$ per period, of order five basis points per year at $r = 1\%$ — below the precision at which the two-bracket class is plausibly distinguishable from a smooth schedule with matched Gini.

8.3 Saez–Zucman proposals through the framework

The Saez–Zucman policy programme (Saez and Zucman, 2019) proposes wealth-tax rates of 2–3% on wealth above \$50M, with explicit appeal to top-share targets. Translated to the present framework via the assessment-fraction rescaling, a 3% nominal top rate corresponds to $r_2^{\text{eff}} \approx 1.65\%$, still below the confinement threshold $km_0 = 2.91\%$ *effective* (equivalently $\approx 5.3\%$ nominal at $\bar{a} = 0.55$) for the Norwegian regime; the 6% proposal sometimes attributed to Saez and Zucman as a maximalist scenario sits just at the threshold. Within the framework, two distinct Saez–Zucman-style designs are visible. A top-share-targeting variant selects high $\alpha_2 \in [3.0, 3.8]$ on the iso-Gini manifold, compressing the upper tail aggressively while leaving the middle bracket relatively flat (Table 1, last rows). A Gini-targeting variant of the same proposal selects moderate $\alpha_2 \approx 1.5$ with a more uniform middle compression. The framework therefore makes precise the informal observation that "the same" top wealth-tax rate implements substantively different distributions depending on the planner's redistributive criterion.

8.4 Where the framework has policy bite

The Norwegian calibration of Section 6.6 locates two regimes for policy use of the framework. Below the wealth-tax-only confinement threshold $r_2^{\text{eff}} > km_0$ — the regime Norway currently occupies — the framework quantifies the marginal contribution of the wealth tax to the joint Pareto exponent and the iso-Gini geometry, but the dominant confinement mechanism sits outside the wealth-tax channel. Above the threshold, the wealth tax itself becomes confining and the criterion-divergence result of Theorem 2 acquires direct policy bite. The framework is therefore most informative for analysing rate-increase proposals on top of the demographic-and-emigration baseline.

8.5 Open directions

Three extensions are most directly suggested by the analysis.

Heterogeneous returns. The derivation throughout treats the regime parameters (m_0, σ) as population-uniform. The flow-versus-stock substitutability of Section 6.6 rests on this assumption, and is precisely the place at which the homogeneous-returns framework reaches its structural limit. Frøseth (2026c) formalises the asymmetry traceable to Domar and Musgrave (1944); Stiglitz (1969): under heterogeneous productive ability z with $\mu(z_H) > \mu(z_L)$, the flow tax compresses the ability-driven drift gap by the factor $\lambda = (1 - \tau_c)(1 - \tau_d) < 1$, while the wealth tax preserves it; a revenue-neutral substitution from flow toward stock taxation therefore widens the drift gap and generates the use-it-or-lose-it efficiency channel of Guvenen et al. (2023). Whether the resulting reallocation is welfare-improving depends on the source of return heterogeneity (productive ability, structural advantage, or winner-take-all dynamics). The analogous design problem — inequality-relevant criteria on a progressive schedule class with ability-stratified (m_0, σ) — is the natural completion of the present paper's arc and the most direct route to a substantively richer flow-versus-stock comparison than the homogeneous-returns calibration of Section 6.6 affords. The framework yields a state-dependent generator that does not collapse to the closed-form OU or piecewise-Laplace structure; numerical treatment is straightforward, closed forms appear unlikely.

Behavioural responses. The framework treats the wealth dynamics as exogenous to the schedule. Behavioural responses — savings, risk-taking, portfolio choice — modify the regime parameters as a function of the schedule itself. Frøseth (2026d,c) characterise the portfolio-choice channel under the (C1)–(C3) class and show that proportional wealth taxes are portfolio-neutral. Whether progressive schedules retain neutrality, and under what restrictions, is open.

Multi-bracket extensions. The two-bracket class is the present paper's focus because it is the form actually implemented. The natural next step is the K-bracket class with $K \geq 3$, where the inequality-relevant policy space grows to $2K-1$ dimensions and a richer family of criterion divergences is expected. Whether the divergence saturates at some finite K or grows without bound is an open quantitative question.

Acknowledgements

The author acknowledges the use of Claude (Anthropic) for assistance with literature review, LaTeX typesetting, mathematical exposition, and editorial refinement, and Lemma (Axiomatic AI) for review and proof checking. All substantive arguments, economic reasoning, and conclusions are the author's own.

A Proofs

This appendix collects the longer technical proofs whose inclusion in the main text would interrupt the narrative flow. The corresponding theorems, propositions, and corollaries are stated in the main text with brief proof sketches; the full arguments follow.

A.1 Proof of Proposition 4 (wealth-tax revenue at steady state)

We compute the two expectations needed. Writing $X = e^Y$ with $Y \sim \mathcal{N}(y_*, \sigma_Y^2)$,

$$\mathbb{E}[X] \;=\; \mathbb{E}[e^Y] \;=\; \exp\Big\{y_* + \tfrac{1}{2}\sigma_Y^2\Big\} \;=\; \exp\left\{y_0 + \frac{km_0 - r_0}{\beta} + \frac{k\sigma^2}{4\beta}\right\}.$$

Using $m_0 = \mu - \sigma^2/2$, the exponent simplifies: $km_0 + k\sigma^2/4 = k\mu - k\sigma^2/2 + k\sigma^2/4 = k(\mu - \sigma^2/4)$, giving the right-hand side of (26).

For $\mathbb{E}[Ye^Y]$ with $Y \sim \mathcal{N}(y_*, \sigma_Y^2)$, the Gaussian moment-generating-function identity yields

$\mathbb{E}[Ye^Y] = \mathbb{E}[e^Y] \cdot (y_* + \sigma_Y^2)$. Therefore,

$$\begin{aligned}
\mathbb{E}[r(X)X] &= r_0\mathbb{E}[X] + \beta\,\mathbb{E}[(Y - y_0)e^Y] \\
&= r_0\mathbb{E}[X] + \beta\mathbb{E}[X]((y_* - y_0) + \sigma_Y^2) \\
&= \mathbb{E}[X](r_0 + \beta(y_* - y_0) + \beta\sigma_Y^2) \\
&= \mathbb{E}[X](km_0 + \beta\sigma_Y^2) \\
&= \mathbb{E}[X](km_0 + k\sigma^2/2) \\
&= k\mu\,\mathbb{E}[X]\,,
\end{aligned}$$

where the fourth line uses $\beta(y_* - y_0) = km_0 - r_0$ from (20) and the fifth uses $\sigma_Y^2 = k\sigma^2/(2\beta)$ from (21). □

A.2 Proof of Proposition 6 (closed-form Gini and top share, single-bracket)

For (38), change variables to $Z = X/X_0$: $\mathbb{E}[Z] = \int_0^1 z \cdot \pi_0 z^{\lambda_-} \mathrm{d}z/X_0 + \int_1^\infty z \cdot \pi_0 z^{\lambda_+} \mathrm{d}z$, both integrals finite under $\alpha > 1$, evaluating to $\pi_0/(\lambda_- + 1) + \pi_0/(\alpha - 1)$. Substituting $\pi_0 = \lambda_-\alpha/(\lambda_- + \alpha)$ and combining gives (38).

For (39), use the standard identity Gini $= (1/\mathbb{E}[X]) \int_0^\infty F(x)(1 - F(x))\,\mathrm{d}x$. The CDF is $F(x) = p_-(x/X_0)^{\lambda_-}$ for $x < X_0$ and $F(x) = 1 - p_+(x/X_0)^{\lambda_+}$ for $x \geq X_0$. Direct integration on each regime, using $\int_0^1 u^a(1 - cu^a)\mathrm{d}u = 1/(a + 1) - c/(2a + 1)$ for the lower regime and the analogous identity (with negative exponents) for the upper, gives the bracketed expression in (39); division by (38) completes.

For (40) at $q \geq p_-$, the quantile $z_q = X_q/X_0$ lies in the upper regime, so $1 - q = p_+ z_q^{\lambda_+}$ inverts to $z_q = (p_+/(1 - q))^{1/\alpha}$. The mean above the quantile, in the Pareto upper regime, satisfies $(1 - q)\,\mathbb{E}[Z \mid Z \geq z_q] = (1 - q)\,\alpha z_q/(\alpha - 1) = \pi_0 z_q^{1-\alpha}/(\alpha - 1)$. Dividing by $\mathbb{E}[Z]$ from (38) (with $X_0 = 1$) and substituting z_q yields (40). □

A.3 Proof of Proposition 7 (single-bracket criterion equivalence)

For S_q, taking logs in (40) and differentiating with respect to α at fixed λ_- gives

$$\frac{\partial \log S_q}{\partial \alpha} = \frac{\log(1 - q)}{\alpha^2} + \frac{1}{\alpha^2} \log\left(\frac{\lambda_- + \alpha}{\lambda_-}\right) - \frac{1}{\alpha(\lambda_- + \alpha)}\,.$$

The first term is strongly negative (since $1 - q$ is small for q near unity), and dominates: numerical evaluation at any $\lambda_- > 0$ and $q \geq 0.5$ confirms the derivative is negative throughout $\alpha > 1$. So S_q is strictly decreasing in α.

For Gini, the strict monotonicity follows from a direct argument. As $\alpha \to 1^+$, the upper Pareto tail has diverging mean and the bulk of the distribution shifts toward arbitrarily high wealth, driving Gini $\to 1$. As $\alpha \to \infty$, the upper regime collapses to a Dirac mass at the threshold and the steady-state distribution shrinks to a finite support, with Gini approaching a finite floor below 1. A direct numerical check across the empirically relevant range $\alpha \in [1.2, 5]$ at any $\lambda_- > 0$ confirms strict monotone decrease.

Both criteria being strictly monotone in α, the target-inversion problems $\text{Gini}(\lambda_-, \alpha) = \text{Gini}^*$ and $S_q(\lambda_-, \alpha) = S_q^*$ each have a unique solution. Equivalent targets pick the same α^*, yielding the same schedule by Corollary 2. $\square$

A.4 Proof of Theorem 2 (criterion divergence, two-bracket)

We show that the gradients are linearly independent at every point where the closed-form expressions (39)–(40) are smooth and finite-mean ($\alpha_2 > 1$).

By Proposition 9 the top share at q above $1 - p_2$ takes the form

$$S_q \;=\; (1-q)^{(\alpha_2-1)/\alpha_2} \cdot H(\lambda_1, \alpha_2, \Delta)\,,$$

where the prefactor depends on α_2 alone — crucially, not on λ_1 — and the remainder H collects the $\mathbb{E}[X]/X_0$, e^{Δ}, and p_2^{1/α_2} contributions, each of which is smooth in $(\lambda_1, \alpha_2, \Delta)$ on the finite-mean regime. The Gini coefficient takes the form $\text{Gini} = G(\lambda_1, \alpha_2, \Delta)$ with no analogous explicit factor in $(1-q)$ and with λ_1-dependence entering through the middle-regime contribution to the Lorenz integral.

To establish linear independence of ∇Gini and ∇S_q, it suffices to exhibit a single point where the 2×2 minor

$$J(\lambda_1, \alpha_2) \;=\; \det\begin{pmatrix} \partial_{\lambda_1}\text{Gini} & \partial_{\alpha_2}\text{Gini} \\ \partial_{\lambda_1} S_q & \partial_{\alpha_2} S_q \end{pmatrix}$$

is non-zero. At the calibrated point $\lambda_0 = 1.53$, $\Delta = 1.5$, $\lambda_1 = 0$, $\alpha_2 = 2.0$ (where $\text{Gini} = 0.505$ and $S_{99} = 0.124$), direct numerical evaluation of the closed forms gives the partial derivatives $\partial_{\lambda_1}\text{Gini} \approx -0.041$, $\partial_{\alpha_2}\text{Gini} \approx -0.090$, $\partial_{\lambda_1} S_{99} \approx -0.011$, $\partial_{\alpha_2} S_{99} \approx -0.123$, yielding $J \approx +4.1 \times 10^{-3}$, distinctly non-zero. The non-vanishing of J on a dense open subset follows from the fact that J, viewed as a function of $(\lambda_1, \alpha_2, \Delta)$, is real-analytic on the finite-mean regime; if J were identically zero on any open subset, real-analytic continuation would force $J \equiv 0$ everywhere, contradicting the observation $J \neq 0$ at the exhibited point. By the implicit function theorem applied at this point, the iso-Gini level set is locally a 2-dimensional C^∞ submanifold of the 3-dimensional policy space, and S_q varies non-trivially on it; the variation is global rather than infinitesimal, with the magnitude given by Table 1. $\square$

A.5 Proof of Proposition 10 (uniform-discount rescaling)

Under uniform assessment, the wealth-tax payment per unit time on raw wealth X is $r(aX) \cdot aX$, where the schedule $r(\cdot)$ takes the assessed-wealth argument aX as its input. The effective rate on raw wealth, defined as the per-unit-wealth flow extracted in tax, is therefore $a \cdot r(aX)$, replacing the no-discount rate $r(X)$ in the post-tax SDE (5). For the smooth schedule $r(A) = r_0 + \beta \log(A/X_0)$,

$$a\, r(aX) \;=\; ar_0 + a\beta \log(aX/X_0) \;=\; a(r_0 + \beta \log a) + a\beta \log(X/X_0)\,,$$

which is again of smooth log-progressive form on raw wealth with effective parameters as claimed. The effective structural ratio $\rho_p^{\text{eff}} = \beta^{\text{eff}}/k = a\beta/k = a\rho_p$ follows; the log-variance $\sigma_Y^2 = k\sigma^2/(2\beta^{\text{eff}}) = \sigma^2/(2a\rho_p)$ follows from Proposition 2. For the bracket schedule $r(A) = r \cdot \mathbf{1}\{A \geq X_0\}$,

$$a\,r(aX) \;=\; ar \cdot \mathbf{1}\{aX \geq X_0\} \;=\; ar \cdot \mathbf{1}\{X \geq X_0/a\}\,,$$

again of single-bracket form on raw wealth, with effective rate ar and threshold X_0/a. The rescaled Pareto exponent $2(ar/k - m_0)/\sigma^2$ follows from Definition 3 with the substitution $r \to ar$. The two-bracket case follows by the same argument applied piecewise, with both thresholds rescaling as $X_i \to X_i/a$ and the ratio X_1/X_0 therefore unchanged. Since the proofs of Proposition 7 and Theorem 2 use only the steady-state shape parameters, the criterion-equivalence and criterion-divergence statements carry through under the rescaling. □

References

Annette Alstadsæter, Niels Johannesen, and Gabriel Zucman. Tax evasion and inequality. *American Economic Review*, 109(6):2073–2103, 2019. doi: 10.1257/aer.20172043.

Roland Bénabou. Tax and education policy in a heterogeneous-agent economy: What levels of redistribution maximize growth and efficiency? *Econometrica*, 70(2):481–517, 2002. doi: 10.1111/1468-0262.00293.

Jean-Philippe Bouchaud and Marc Mézard. Wealth condensation in a simple model of economy. *Physica A*, 282(3–4):536–545, 2000. doi: 10.1016/S0378-4371(00)00205-3.

Evsey D. Domar and Richard A. Musgrave. Proportional income taxation and risk-taking. *Quarterly Journal of Economics*, 58(3):388–422, 1944. doi: 10.2307/1882847.

Anders G. Frøseth. Tax migration as social contagion: A tipping-point model with application to the Scandinavian wealth tax debate. 2026a. arXiv:2607.03868 [physics.soc-ph].

Anders G. Frøseth. Minimum-distortion wealth taxation, i: Information-theoretic versus transport-geometric optimality on the proportional class. Working paper, 2026b.

Anders G. Frøseth. Extensions to the wealth tax neutrality framework. 2026c. arXiv:2603.05277 [physics.soc-ph].

Anders G. Frøseth. Flow taxes, stock taxes, and portfolio choice: A generalised neutrality result. 2026d. arXiv:2603.15974 [physics.soc-ph].

Anders G. Frøseth. Asset returns, portfolio choice, and proportional wealth taxation. 2026e. arXiv:2603.05264 [physics.soc-ph].

Anders G. Frøseth. From gravity to confinement: Wealth redistribution as optimal drift design in the Fokker–Planck framework. 2026f. arXiv:2607.06153 [physics.soc-ph].

Anders G. Frøseth. Wealth taxation as a drift modification: A Fokker–Planck approach to tax neutrality. 2026g. arXiv:2603.05283 [physics.soc-ph].

Anders G. Frøseth. Wasserstein geometry of wealth taxation: Rigid motions of the log-wealth distribution and the optimality of progressive wealth taxes. Working paper, 2026h.

Xavier Gabaix. Zipf's law for cities: An explanation. *Quarterly Journal of Economics*, 114(3): 739–767, 1999.

Xavier Gabaix, Jean-Michel Lasry, Pierre-Louis Lions, and Benjamin Moll. The dynamics of inequality. *Econometrica*, 84(6):2071–2111, 2016.

Fatih Guvenen, Gueorgui Kambourov, Burhanettin Kuruscu, Sergio Ocampo-Diaz, and Daphne Chen. Use it or lose it: Efficiency and redistributional effects of wealth taxation. *The Quarterly Journal of Economics*, 138(2):835–894, 2023. doi: 10.1093/qje/qjac047.

Jonathan Heathcote, Kjetil Storesletten, and Giovanni L. Violante. Optimal tax progressivity: An analytical framework. *The Quarterly Journal of Economics*, 132(4):1693–1754, 2017. doi: 10.1093/qje/qjx018.

NOU 2022:20. Et helhetlig skattesystem. Technical report, Norges offentlige utredninger, Oslo, 2022.

Thomas Piketty. *Capital in the Twenty-First Century*. Harvard University Press, Cambridge, MA, 2014.

Thomas Piketty and Gabriel Zucman. Capital is back: Wealth-income ratios in rich countries 1700–2010. *Quarterly Journal of Economics*, 129(3):1255–1310, 2014.

Emmanuel Saez and Gabriel Zucman. Progressive wealth taxation. *Brookings Papers on Economic Activity*, 2019(Fall):437–533, 2019. doi: 10.1353/eca.2019.0017.

Peter Birch Sørensen. Neutral taxation of shareholder income. *International Tax and Public Finance*, 12(6):777–801, 2005. doi: 10.1007/s10797-005-3010-8.

Joseph E. Stiglitz. The effects of income, wealth, and capital gains taxation on risk-taking. *Quarterly Journal of Economics*, 83(2):263–283, 1969.

Joseph E. Stiglitz. Equilibrium wage distributions. *The Economic Journal*, 95(379):595–618, 1985. doi: 10.2307/2233029. JSTOR stable URL: https://www.jstor.org/stable/2233029.